

Presuppositions and Default Reasoning: A Study in Lexical Pragmatics

Robert E. Mercer

*Department of Computer Science
Middlesex College
University of Western Ontario
London, Ontario, CANADA
N6A 5B7
mercer@csd.uwo.ca*

Abstract

Explaining how the meaning of words relate to the meaning of the utterance in which they are used is of utmost importance. The most common approaches view the meaning of an utterance as a composition of the meanings of its parts, which of course include the words used to construct the utterance. This approach is successful for entailments. However, similar approaches to explain the presuppositional behaviour of utterances have for the most part failed. In this paper we describe the application of Default Logic to the representation and the generation of natural language presuppositions. The view is taken that the presuppositions of an utterance are conjectures made by the hearer based upon the assumption that the speaker is following Grice's maxims of cooperative conversation. These conjectures represent information implicitly contained in the utterance which cannot be generated by classical techniques. The compositional framework is maintained. The difference is that functional units rather than predetermined semantic units are inherited by the meaning structure. The function's meaning changes depending on the contents of the meaning structure. Hence, we view the study of these functional units as *lexical pragmatics* rather than lexical semantics. Default Logic is one formal method for performing default reasoning in the area of Artificial Intelligence called Knowledge Representation. Default reasoning attempts to fill with conjectures the gaps left by classical forms of reasoning. We suggest that the use of non-classical inferencing techniques such as default reasoning will prove fruitful in the realm of lexical reasoning.

1 Introduction

Lexical semantics deals with various aspects of the meanings of words. One such aspect is entailment. The standard meaning of the logical notion of entailment can be naturally extended to natural language sentences (see [Cruse, 1986], for example). We can abstract away from the linguistic setting and formalize this linguistic entailment relation, by considering *meaning postulates* and the logical entailment relation (usually as some form of deduction).

Here, we are interested in the kinds of meanings that we can give to words in various contexts which puts the meaning postulate within the scope of a negation. We again wish to abstract away from the linguistic setting and formalize this process. We accomplish this using Default Logic. Exactly how this formal method transfers back to the practical linguistic setting is left for future investigation.

We have coined the expression *lexical pragmatics* to emphasize the significant differences between the semantic and pragmatic aspects of a lexical item's meaning. Whereas

the semantics of a lexical item remains constant in all contexts, and hence can be pre-computed, the pragmatics of a lexical item is determined by the context in which it is found. The pragmatics of a lexical item must then be represented as a function which is computed in context. The use of default rules shows precisely what the pragmatic notion needs to be.

2 Background

What is a natural language presupposition? Being implied by a simple natural language sentence and the natural (or preferred) interpretation of its simple negation is the primary quality that qualifies an inference as a presupposition. This evaluation of inferences is called the *negation test*. Presuppositions are generated from lexical and syntactic contexts. Those contexts which pass the negation test can be termed *presuppositional triggers*. The presuppositional trigger that is of concern here is the definition of words. In (1) the truth of the affirmative *a*-sentence always implies the truth of the *c*-sentence, and the truth of the negative *b*-sentence normally implies the truth of the *c*-sentence.

- (1) a. My cousin is a bachelor.
b. My cousin is not a bachelor.
c. My cousin is a male adult.

The natural language presupposition has been a well-studied linguistic phenomenon. Because it has received so much attention, it has naturally been given a variety of definitions. So, before presenting the technical details of our work, it should be clarified that the point of view of Gazdar [1979] is assumed. In particular, presuppositions are considered as part of the information conveyed by an utterance (*not* appropriacy conditions for making an utterance) given that the utterance was made cooperatively in the sense of Grice's maxims of cooperative conversation [Grice, 1975]. So, in the case of (1b) the view would be that what is being communicated is that '*the speaker knows¹ that my cousin is a married male adult.*', not that it is inappropriate to utter the sentence in contexts which don't allow for the truth of '*My cousin is a male adult.*'. Although the speaker-hearer view of communication is implicit in Gazdar's approach, we have made it explicit in our communication model. This has been done because we are interested in what information is needed and what forms of inference are required to enable a hearer to infer the correct presuppositions from an utterance.

Presuppositions have the quality that there are linguistically natural means for indicating that a simple negation is not to be interpreted normally, that is, some form of "presupposition cancellation" is to be performed. (What is meant by the terminology in scare quotes also differs from method to method.) It is this latter problem, how to produce only the correct presuppositions of sentences, that has received most of the attention in the literature. It is called the *projection problem*. From a lexical point of view, the projection problem can be interpreted as how do the meanings of the lexical items contribute to the meaning of the utterance.

Examples are given in (2). It is important to note that two methods are used to indicate that the negation is to be interpreted in an unnatural manner. (2a) illustrates

¹See [Gazdar, 1979] for a discussion of the motivation for the epistemic qualification of the proposition. Gazdar uses the usual omniscient interpretation of the epistemic operator, Know. This causes some problems which are not important here.

that one method is to make the normal (presuppositional) inference inconsistent. (2b) indicates a second method. This sentence is actually a disguised form for '*My cousin is a bachelor and not a spinster (case 1) or my cousin is a spinster and not a bachelor (case 2)*'. The speaker cannot commit himself to the natural interpretation of the negations because each of the cases commits the speaker to deny the presuppositions required by the natural interpretation. Details are given in the next section.

- (2) a. My cousin is not a bachelor. He is only three years old.
- b. My cousin is a bachelor or a spinster.

Note that the sentences in (3) do interpret '*bachelor*' in the normal way under normal circumstances.

- (3) a. My cousin is a bachelor or Mary will be unhappy.
- b. If your cousin is a bachelor then he will be able to attend the party.

3 Presuppositions and Default Logic

Throughout our work the projection problem has been approached from a knowledge representation point of view. A basic assumption has been that a hearer's interpretation of a natural language utterance should include the inferences that can be generated from three sources: the sentence uttered, knowledge about the world, and knowledge about language use.

For the purposes of this paper the most important knowledge source is the lexicon. It is here that some of the knowledge about language use is stored (for our purposes the definitions of words). Other linguistic knowledge, such as Grice's maxims, has been placed in the general inference procedure. What information is to be stored in the lexicon and in what form has been the focus of this study into the phenomenon called presupposition.

Previously, the standard "solution" to the projection problem has been one of three kinds: a set of rules that take the presuppositions of the clauses and remove the undesirable presuppositions as the sentence meaning is being composed [Karttunen, 1973, Karttunen, 1974, Karttunen and Peters, 1979, Soames, 1979]; a set of rules, invoked after the sentence has been fully interpreted, that cancel the unwanted presuppositions from a complete set of potential presuppositions [Gazdar, 1979]; or a set of rules that embody both of these methods [Soames, 1982].

In the work discussed below, we have rejected these solutions and have replaced them with a method that views presupposing as a form of inference. Before giving details of the method, it would be appropriate to point out the philosophical differences between this new method and the previous ones, indicating both the need for default reasoning and the benefits that are obtained.

Previous methods and the new one agree that the information concerning the presuppositions associated with words is located in the lexicon. Where they differ is what form this information takes. This representation has important ramifications on the kinds of knowledge needed outside the lexicon. More specifically, previous methods put extensional knowledge in the lexicon, that is, the knowledge of a particular lexical item is designated as presuppositional and non-presuppositional (or what has been termed "what is said"). So in the case of '*bachelor*', '*unmarried*' would be designated as non-presuppositional; '*male*' and '*adult*', presuppositional. What this precomputing of

presuppositions requires is that knowledge be contained in a projection method which can eradicate the unwanted presuppositions in sentences like (2). [Karttunen, 1973, Karttunen, 1974], [Karttunen and Peters, 1979], and [Soames, 1982] exemplify this distribution of knowledge. [Gazdar, 1979] did not have to resort to such a distribution of knowledge, however this method succumbs to other problems. The new method has evolved naturally from Gazdar's ideas but does not succumb to the same problems.

The new method differs in the following important respect: the information about presuppositions is stored in the lexicon as intensional information, that is, in a pre-computed form. As an utterance is processed to derive its meaning, the intensional information for all of the presuppositional triggers is incorporated into the structure which represents the meaning of the utterance. Later, we will see that an appropriate structure is a default theory. Since the intensional information is then computed in the appropriate context (the structure representing the utterance's meaning), the presuppositional information is correctly generated. By formulating the problem of presupposition generation in this way, the importance of the complete context for the correct generation of an utterance's presuppositions dramatically exposes the pragmatic nature of presuppositions. Hence, the kind of information that is stored in the lexicon and the kind of processing used to connect the information to the structure representing the meaning of the utterance tends to suggest that it should be under the heading of *lexical pragmatics* rather than *lexical semantics*.

This new method can be viewed as a reasonable extension to Carnap's *meaning postulates* [Carnap, 1956]. The meaning postulate is extended to include the intensional information, which can be given a precise meaning. The extensional approach on the other hand cannot be used to extend the idea of a meaning postulate since the notion of the universality of the meaning postulate is lost.

Another important repercussion is that the previous methods are *ad hoc*, that is, they are not based on fundamental principles that have their foundations outside the particular method. On the other hand the default method can actually provide a better understanding of the phenomenon and leads to a definition based on foundations which are not intimately connected to the phenomenon and which are not complicated by operational features.

Other benefits are obvious. We are dealing with a very large task when we confront the natural language understanding problem. We don't need to compound the difficulty of this task with purported solutions that lead to knowledge needlessly distributed amongst various parts of the solution. The default reasoning method has succeeded in making what seems to be a most natural break, that is, the knowledge associated with the lexical entries is located in the lexicon and the knowledge concerning the basic inference mechanism is located in the inference engine.

3.1 Details of the Representation

The work on presuppositions and default reasoning began with [Mercer and Reiter, 1982] reporting on its embryonic stages. Here we looked at representing presuppositions as default rules in Reiter's [1980] Default Logic. More recently, we have viewed this method as a natural way to extend Carnap's notion of a meaning postulate. This extension differs from the original in two significant ways. Firstly, we make the change from a one-way material implication in the meaning postulate to a logical equivalence. So, we are not only accepting that the right hand side of (4) are necessary elements of the definition

of 'bachelor', we are also accepting that these are sufficient, as well. Much debate has focussed on this issue. We agree that it is probably impossible to arrive at a necessary and sufficient set of conditions for the definition of a lexical item. However, this is not a theory about the correct definition of lexical items. Rather, it is an agent-centred method about how an agent uses his knowledge to derive other facts. Hence, we do take as given that at any moment an agent does have necessary and sufficient conditions defining a lexical item. This definition may be subject to change. How an agent discovers that his definition needs modification, and how this is performed is not part of our immediate concern. If (4) were a one-way material implication, the contrapositive of (4) would say that the negation of any of the three right hand side predicates is sufficient to negate 'bachelor'. However, it does not say that the negation of 'bachelor' necessarily requires that the negation of one of the right hand side predicates is the cause. What having necessary and sufficient conditions allows an agent to do is determine that one of the right hand side items in (5) is causing the negation of 'not a bachelor' although it cannot be determined which one.

$$(4) \quad \forall x. BACHELOR(x) \equiv MALE(x) \wedge ADULT(x) \wedge \neg MARRIED(x)$$

$$(5) \quad \forall x. \neg BACHELOR(x) \equiv \neg MALE(x) \vee \neg ADULT(x) \vee MARRIED(x)$$

Secondly, we add more information about the lexical item in question. This extra information concerns the presupposed² parts when the lexical item occurs in the scope of a negation. This extra information takes the form of default rules and implicitly adds a new form of inference, as well. So for example, we add to the representation of 'bachelor' the two default rules (6) and (7).

$$(6) \quad \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : MALE(x)}{MALE(x)}$$

$$(7) \quad \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : ADULT(x)}{ADULT(x)}$$

A more detailed discussion of Default Logic is given in the next section. For the present, an interpretation of these two default rules is sufficient. (6) is to be interpreted as stating that if the non-bachelorhood of x is provable from the default theory (x will be substituted appropriately throughout the rule with the appropriate constant) and if the linguistic form (that is, the sentence) predicates bachelorhood about x , then if it is consistent in the default theory to infer that x is male then infer that x is male. (7) says the same thing about x and the property of being an adult. Default rules will be used to determine which of the right hand side elements is causing the negation. So, under normal circumstances (6) and (7) will generate $MALE(x)$ and $ADULT(x)$ which together with (5) gives the result that $MARRIED(x)$.

²Earlier work has called these the noncriterial parts. The one part negated by the preferred interpretation was called the criterial part. This terminology was used because we were concerned with describing the negated and unnegated parts of the preferred interpretation without referring to *presuppositions*. Since this terminology differs from the use of *criterial* in [Cruse, 1986], say, (each of *male*, *adult*, and *unmarried* would be criterial), and since it is not so critical here to be so pedantic about the term *presupposition*, we will use this latter terminology.

3.2 Details of the Generation Method

The study of the method using Default Logic as a replacement for the projection methods is the main theme of [Mercer, 1987]. A summary of this work appears in [Mercer, 1988b]. An important contribution is the connection made between presuppositions and the context, in particular Gazdar's formalization of Grice's maxims. More precisely, default theories corresponding to a well-defined set of cases which can be derived from the hearer's knowledge base (KB_H) and the utterance (represented as $\mathcal{G}(u)$) to convey Gazdar's view of the Gricean maxims) generate the appropriate presuppositions for sentences such as those found in (2) and (3). The precise details for computing the set of cases (which can be found in [Mercer, 1987, Mercer, 1988b]) are not important for this paper. Instead, some examples will be provided in order that an understanding of how the Default Logic method generates the appropriate presuppositions for a variety of sentences. The default theories that correspond to the cases mentioned above will be given without any explanation. Interested readers are encouraged to seek further explanation in the references given above. More recently, [Mercer, 1988a, Mercer, 1990] have solved problems left open in [Mercer, 1987], specifically, how negation interacts with sentential adverbs in the context of conditionals, and a generalization of Gazdar's view of clausal implicatures. The default method now captures all of the known problems for any of the previous methods plus the problem of disjunctive presuppositions which none of the previous methods could handle.

Before embarking on a study of the application of Default Logic to the problem of natural language presuppositions, a brief introduction to Default Logic is in order. The inference rules of classical logic produce *safe* deductions. Deductions are safe if they are, in some sense, already known. Examples of this type of inference include: deducing that a particular individual has a property from the knowledge that all individuals have that property; deducing the truth of two conjuncts given the truth of the conjunction; and deducing the consequent of a material implication given the truth of the antecedent.

Default Logic, on the other hand, allows (in a principled manner) the conjecturing of new information in the face of incomplete information. Remembering the discussion surrounding (5), the representation of a negated lexical item does not provide complete information, and since a more complete interpretation of the negated lexical item is required, conjecturing of the extra information is necessary.

A *default rule* is a rule of inference as shown in (8),

$$(8) \quad \frac{\alpha(\vec{x}) : \beta_1(\vec{x}), \dots, \beta_n(\vec{x})}{w(\vec{x})}$$

where $\alpha(\vec{x})$, $\beta(\vec{x})$, $w(\vec{x})$ are all first order formulae whose free variables are among those of $\vec{x} = x_1, \dots, x_m$. Intuitively, a default rule can be interpreted as *for all individuals x_1, \dots, x_m , if $\alpha(\vec{x})$ is believed and if $\beta(\vec{x})$ is consistent with our beliefs, then $w(\vec{x})$ may be believed*. The α , β_i 's, and w are referred to as the *prerequisite*, *justifications*, and *consequent* of the default rule, respectively.

Normal defaults are default rules with the properties $n = 1$ and $\beta_1(\vec{x}) = w(\vec{x})$. Since all the default rules used throughout this paper are normal, discussion is restricted to this class. In addition to this structural property, normal defaults have other important properties which are discussed below. *Closed defaults* are default rules that contain no free variables in any of the first-order formulae. *Open defaults* are default rules that are not closed. An open default is meant to represent the set of closed defaults obtained

by replacing all its free variables by ground terms. *Default rule schemata* are meta-default rules that represent the set of open default rules obtained by replacing all the meta-variables with the appropriate substitution instance (for example, predicates, or formulae). A *default theory*, Δ , is composed of a set of first-order formulae W , and a set of default rules, D . The default rules can be viewed as *extending* the first-order formulae with the consequents of the default rules. An *extension*, E , of a closed default theory is a constructable fixed point having the properties listed below.

1. $W \subseteq E$
2. if $E \vdash a$, then $a \in E$ (E is logically closed, that is, $Th(E) = E$, where Th is a fixed point operator defined by \vdash)
3. for each default, $\frac{\alpha(\vec{x}) : \beta(\vec{x})}{\beta(\vec{x})} \in D$, if $\alpha(\vec{x}) \in E$, and $\neg\beta(\vec{x}) \notin E$, then $\beta(\vec{x}) \in E$

Normal default theories always have an extension. Multiple extensions are produced when two defaults conflict. Normal defaults conflict when their consequents cannot simultaneously be in the same extension. These multiple extensions are orthogonal, that is, pairwise, they differ on the truth value of at least one formula. The third point in the definition of E is very important. Limiting the discussion to the effect that this definition has on the application discussed in the remainder of this paper, since the first-order part of the default theory and all its entailments are in E , if a conflict occurs between a default and anything in the logical closure of W the consequent of the default rule is in no extension of Δ .

Although Default Logic has a proof theory, it will not be described here. Rather all of the discussion regarding presuppositions will use the definition of extensions just given. We are able to do this since the logic has a kind of soundness and completeness result. That is, the proof theory cannot generate any formulae that is not in some extension, and if a formulae is in an extension, then the proof theory can generate it.

For purposes of all the following examples, the definition of 'bachelor' is represented by the first order sentence (4). Then the negation of 'bachelor' would be represented by (5). Similarly the definitions for 'spinster', and its negation, are given by (9) and (10), respectively.

$$(9) \quad \forall x. SPINSTER(x) \equiv FEMALE(x) \wedge ADULT(x) \wedge \neg MARRIED(x)$$

$$(10) \quad \forall x. \neg SPINSTER(x) \equiv \neg FEMALE(x) \vee \neg ADULT(x) \vee MARRIED(x)$$

The pragmatic rule for generating presuppositions from all lexical items can be captured as a default rule *schema* (11), where PRE represents information about the presuppositional features of the negated lexeme. The knowledge of presuppositional features could then be part of the lexicon and could be represented as (12).

$$(11) \quad \frac{\neg P(x) \wedge LF(P, x) : P_1(x)}{P_1(x)}, \text{ if } PRE(P, P_1) \text{ is in the lexicon}$$

$$(12) \quad \begin{array}{l} PRE(BACHELOR, MALE) \\ PRE(BACHELOR, ADULT) \\ PRE(SPINSTER, FEMALE) \\ PRE(SPINSTER, ADULT) \end{array}$$

This method is just an economical form of representation. Since (11) is a schema it really represents the set of its instances. These instances can be obtained as needed by appropriately substituting for P and P_1 throughout the schema. For the case of ‘*bachelor*’ we have already seen the instances as (6) and (7). The instances for ‘*spinster*’, are similar. All of the examples below will use the already substituted instances rather than the schema. Before turning to the examples, we first need a more precise definition of presuppositions based on defaults.

3.2.1 A Definition of Presuppositions

In addition to the broad coverage, the default method provides a definition of presupposition not forthcoming from the previous methods. Before giving the definition of presuppositions, we require the prior definition of the consequents of the default theory.

Definition 1 *If D is any set of defaults then*

$$CONSEQUENTS(D) = \left\{ w(\vec{x}) \left| \frac{\alpha(\vec{x}) : \beta_1(\vec{x}), \dots, \beta_n(\vec{x})}{w(\vec{x})} \in D \right. \right\}$$

For purposes of the next definition, the only defaults in KB_H , are the presupposition-generating defaults found in the speaker’s lexicon.

Definition 2 *Let u be a sentence uttered by a speaker, S , in accordance with Grice’s Maxims of Cooperative Conversation. Let KB_H be the hearer’s knowledge base before the utterance, and let the default theories $\Delta_{u_{Case1}}, \dots, \Delta_{u_{Case n}}$ be the first order cases of the theory $KB_H \cup \{G(u)\}$. A sentence α is a presupposition of u with respect to KB_H if and only if*

- (i) α is in all extensions of $\Delta_{u_{Case i}}$, for $i = 1, \dots, n$,
- (ii) α is in the deductive closure of $CONSEQUENTS\{D\}$
- (iii) α is not an entailment of $KB_H \cup \{G(u)\}$
- (iv) α is not inferable from KB_H

3.2.2 Example 1 — Simple Negation

In the situation in which ‘*My cousin is not a bachelor.*’ has been uttered, and the logical form created maps ‘*my cousin*’ onto the constant c_1 , the representation of the sentence would be that given in (13). However, included in the hearer’s lexicon is the information provided by the extended meaning postulate for ‘*bachelor*’ represented by (14)–(16). These four items taken together provide the default theory that represents the meaning of this utterance.

$$(13) \neg BACHELOR(c_1)$$

$$(14) \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : MALE(x)}{MALE(x)}$$

$$(15) \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : ADULT(x)}{ADULT(x)}$$

$$(16) \forall x. \neg BACHELOR(x) \equiv \neg MALE(x) \vee \neg ADULT(x) \vee MARRIED(x)$$

It can easily be seen that all extensions of the default theory consisting of (13)–(16) (there is only one) contains $MALE(c_1)$ and $ADULT(c_1)$. By the definition given in the previous section, these are the lexical presuppositions of the utterance derived from ‘*bachelor*’.

3.2.3 Example 2 — Negation with Qualification

In the situation in which ‘*My cousin is not a bachelor, because my cousin is just three years old.*’ has been uttered, and the logical form created maps ‘*my cousin*’ onto the constant c_1 , the representation of the sentence would be that given in (17). Again, included in the hearer’s lexicon is the information provided by the extended meaning postulate for ‘*bachelor*’ represented by (18)–(20). These four items taken together provide the default theory that represents the meaning of this utterance.

$$(17) \neg BACHELOR(c_1) \wedge \neg ADULT(c_1)$$

$$(18) \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : MALE(x)}{MALE(x)}$$

$$(19) \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : ADULT(x)}{ADULT(x)}$$

$$(20) \forall x. \neg BACHELOR(x) \equiv \neg MALE(x) \vee \neg ADULT(x) \vee MARRIED(x)$$

It can be easily seen that all extensions of the default theory consisting of (13)–(16) (there is only one) contains $MALE(c_1)$ but does not contain $ADULT(c_1)$. The reason is that $\neg ADULT(c_1)$ is an entailment of this default theory which defeats the justification of the default rule (19) thereby not allowing its consequent into the extension. Hence $MALE(c_1)$ is a lexical presupposition of the utterance, but $ADULT(c_1)$ is not.

3.2.4 Example 3 — Multiple Cases with No Cancellation

In the situation in which ‘*My cousin is a bachelor or Mary will be unhappy.*’ has been uttered, and the logical form created maps ‘*my cousin*’ onto the constant c_1 , and ‘*Mary*’ onto the constant m_1 , the representation of the sentence would be that given by the two cases (21) and (22). As stated earlier, the computation of these cases is not given here. However, it should be obvious that these two cases are the two disjuncts of the exclusive or representation of the utterance.

$$(21) BACHELOR(c_1) \wedge \neg UNHAPPY(m_1)$$

$$(22) \neg BACHELOR(c_1) \wedge UNHAPPY(m_1)$$

$$(23) \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : MALE(x)}{MALE(x)}$$

$$(24) \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : ADULT(x)}{ADULT(x)}$$

$$(25) \forall x. BACHELOR(x) \equiv MALE(x) \wedge ADULT(x) \wedge \neg MARRIED(x)$$

$$(26) \quad \forall x. \neg BACHELOR(x) \equiv \neg MALE(x) \vee \neg ADULT(x) \vee MARRIED(x)$$

Together with the lexical information that is available for 'bachelor', which is given in (23)–(26) we obtain $MALE(c_1)$ and $ADULT(c_1)$ in all extensions of all cases. So, these are presuppositions of the utterance.

3.2.5 Example 4 — Multiple Cases with Cancellation

In the situation in which 'My cousin is a bachelor or a spinster.' has been uttered, and the logical form created maps 'my cousin' onto the constant c_1 , the representation of the sentence would be that given by the two cases (27) and (28).

$$(27) \quad BACHELOR(c_1) \wedge \neg SPINSTER(c_1)$$

$$(28) \quad \neg BACHELOR(c_1) \wedge SPINSTER(c_1)$$

$$(29) \quad \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : MALE(x)}{MALE(x)}$$

$$(30) \quad \frac{\neg BACHELOR(x) \wedge LF(BACHELOR, x) : ADULT(x)}{ADULT(x)}$$

$$(31) \quad \forall x. BACHELOR(x) \equiv MALE(x) \wedge ADULT(x) \wedge \neg MARRIED(x)$$

$$(32) \quad \forall x. \neg BACHELOR(x) \equiv \neg MALE(x) \vee \neg ADULT(x) \vee MARRIED(x)$$

$$(33) \quad \frac{\neg SPINSTER(x) \wedge LF(SPINSTER, x) : FEMALE(x)}{FEMALE(x)}$$

$$(34) \quad \frac{\neg SPINSTER(x) \wedge LF(SPINSTER, x) : ADULT(x)}{ADULT(x)}$$

$$(35) \quad \forall x. SPINSTER(x) \equiv FEMALE(x) \wedge ADULT(x) \wedge \neg MARRIED(x)$$

$$(36) \quad \forall x. \neg SPINSTER(x) \equiv \neg FEMALE(x) \vee \neg ADULT(x) \vee MARRIED(x)$$

$$(37) \quad \forall x. MALE(x) \supset \neg FEMALE(x)$$

The first case, represented by (27), together with the lexical information represented by (29)–(37) form a default theory which allows $MALE(c_1)$ and $ADULT(c_1)$ to be derived as entailments. Consequently, no extension of this theory contains $FEMALE(c_1)$, the consequent of the default rule (33), since the justification is defeated by $MALE(c_1)$. $ADULT(c_1)$ can also be derived using the default rule (34).

Similarly, in the second case, represented by (28), together with the lexical information represented by (29)–(37) form a default theory which allows $FEMALE(c_1)$ and $ADULT(c_1)$ to be derived as entailments. Since the justification of default rule (29) is defeated by $FEMALE(c_1)$, no extension of this theory contains $MALE(c_1)$. $ADULT(c_1)$ can also be derived using the default rule (30).

Since not all extensions contain $MALE(c_1)$ and not all contain $FEMALE(c_1)$, neither of these formulae are presuppositions of the utterance. Although $ADULT(c_1)$ is in all extensions of all the cases it is not a presupposition of the utterance since it is entailed by it.

3.3 Lexical Pragmatics

As mentioned earlier, we consider that this enterprise falls under the heading of lexical pragmatics. Now that a short description of this study of presuppositions has been given, it is appropriate to make more precise what we mean by the term, lexical pragmatics.

The key difference between semantics and pragmatics has always been the role of context. It is evident from the previous discussion that for natural language presuppositions the role of context is paramount. There are explicit negations which produce different outcomes depending on the (non)existence of further qualifications. Sentences which have exactly the same syntactic structure, '*A or B*,' have different presuppositional behaviours depending upon the content of '*A*' and '*B*'. Although the role of context is fundamental, here we would like to focus instead on the functional units, the default rules, in order to make more direct connections back to the source of these ideas: knowledge representation and default reasoning in Artificial Intelligence.

Once the semantic meaning of '*bachelor*' has been agreed upon, (38) say, this meaning remains unchanged no matter what context it is put in. This property is referred to as monotonicity. Monotonicity requires the behaviour of the inference mechanisms to be localized. All classical inference mechanisms are monotonic. So, for instance, the maleness of a person is always the case given that the person is a bachelor. Consequently, these attributes can be precomputed and associated with the lexical item '*bachelor*'.

$$(38) \text{ MALE}(x) \wedge \text{ADULT}(x) \wedge \neg \text{MARRIED}(x)$$

On the other hand the meaning of '*not a bachelor*' has two components: the semantic component, (39), and the pragmatic component, which is most commonly defined in terms of what is normally the case. Again, the semantic component is unchanging. It states that at least one of the disjuncts is true whenever the person is '*not a bachelor*'. On the pragmatic side, our approach is to reject the normal case only if something contradicts it. This non-monotonic form of reasoning is demonstrated by Default Logic. The default rules make explicit what it is that can defeat the generation of the consequence of that rule. These rules are intensional objects whose meanings are made explicit only in context. However, these rules can still be attached directly to the lexical item in exactly the same manner as the precomputed semantic components.

$$(39) \neg \text{MALE}(x) \vee \neg \text{ADULT}(x) \vee \text{MARRIED}(x)$$

The contextual nature of the presupposition inference as captured by default reasoning indicates that one difference between semantics and pragmatics revolves around the monotonic/non-monotonic character of semantic and pragmatic inferences. This view adds a new dimension to pragmatics. We believe that it is this view that has the most to offer the lexicon.

4 What Remains

How to apply this method in a practical implementation still needs to be investigated. Some of the remaining problems are in the domain of research in Default Logic. Default Logic is not even semi-decidable in the first order case. In the propositional case it is intractable even for syntactically very restricted forms of theories. One issue which has recently become a focus of attention in the default reasoning area is the notion of scope,

that is, when deciding whether the default succeeds, what part of the theory should be investigated for the potentially defeating information. [Etherington *et al.*, to appear] The interest in scope is twofold: to make default reasoning more tractable, and the commonly held belief that defaults in general are somewhat localized in their influence. For our purposes here, the sentence presents a very natural scope. It is probably not coincidental that most of the examples given above deal with intrasentential competition among presuppositional triggers and entailments.

Although this view of scoping can be used to overcome a variety of problems, some of the very (in)famous examples of presupposition deal with noun phrases. Here, we come in contact with referring phrases, which again force us into potentially considering every object in our domain as a possible argument to defeat the default assumption. However, we believe that a proper reading of Grice and an application of Gazdar's epistemic view of communication can provide the needed relief from this problem (although the details have yet to be worked out). In particular, if the speaker uses a particular referring expression in the utterance, then the speaker is suggesting that it is under *this* description that the information is stored. Thus the speaker is also communicating that he knows the information being communicated is related to this descriptor rather than some other descriptor. Our task of deciding what the speaker is communicating is therefore limited to this descriptor. (This of course requires an epistemic operator that is not omniscient.) This limitation assumes that there is no commonly known synonym for the referring expression (which the speaker would not need to include in the utterance). If common synonyms exist, then we regard this as a problem of knowledge representation, that is, how do we deal with these localized equalities without recourse to a promiscuous equality relation.

Earlier we suggested that we have been interested in formalizing the notion of presupposition. Inherent in this task is the use of a formal apparatus, in this case, Default Logic. The application of these ideas to the practical situation of lexicon design and use will require some more effort. For example, the lexical presuppositions that have been discussed are attached to particular lexical items. So, for instance, the presuppositions of '*not a bachelor*' do not hold for the logically equivalent expression '*not an unmarried adult male*' nor does '*female*' which logically entails '*not a bachelor*' allow the default assumption of '*adult*' to hold. In order to prevent these unwanted inferences, we have used guards on the default rules (the *LF* predicate) to limit them to appropriate contexts. Another method would be to allow defaults to enter the inferencing process for limited periods of time rather than being constantly available. How this method compares with the standard view of default rules in a Default Logic remains to be investigated. Once these questions are answered, the attachment of the defaults to lexical items in the lexicon and their proper use in computing the meaning of an utterance can be achieved.

Lastly, it is of some importance to see what questions can be asked given this view of lexical features. Figure 1 represents a small portion of an *ISA* hierarchy of some hearer, *H*. The arcs have been labelled with features and the nodes have been labelled by the term which is associated with the set of features defined by the path from the root node. The deciding criterion for a negated lexeme's meaning is associated with the last feature which labels the last arcs in the paths from the root node to the negated and unnegated lexical nodes. In this instance the non-presuppositional feature is *MARRIED* — \neg *MARRIED*. The emboldened path in Figure 1 represents the preferred interpretation of '*not a bachelor*', that is, '*a married adult male*'. Given this simplistic taxonomic structure two questions arise. Can the lexicon be organized in such a way that the

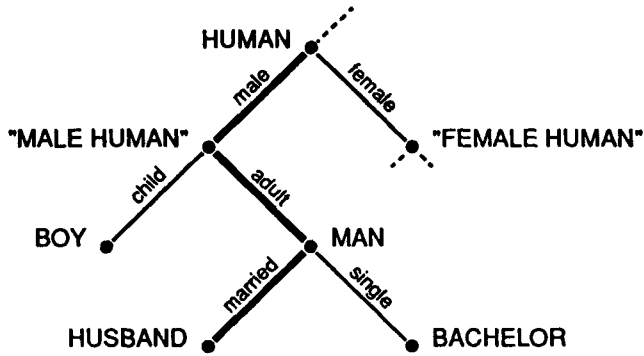


Figure 1: Portion of *H*'s ISA hierarchy.

presuppositional and non-presuppositional features are given implicitly by the structure rather than having to be explicitly represented? Negation would normally modify only the last arc. Given that this type of representation can be achieved are there any interesting answers that arise from the structure, that is, are presuppositions given by the way we structure our linguistic (world) knowledge?

5 Conclusion

We have investigated a formal approach to the definition of natural language presuppositions. It has involved defining the presuppositions of lexical items as default rules together with the use of Default Logic proof theory in appropriate contexts to generate the presuppositions of utterances containing the lexical items that generate presuppositions.

This is but one use of defaults in the lexicon. Other uses are fairly obvious — and they are obvious in a sense that makes the boundary between the lexicon and classical knowledge representation very fuzzy. For example, the notion of canonical traits is the classical example given to support the use of default reasoning in knowledge representation: Tweety is a bird, therefore it (probably) flies. As well, what happens when a standard concept is modified, for example, dog *vs* toy dog. The issue of what gets negated is very similar in nature to the problem of presuppositions: one item in the concept *dog* is negated everything else (by default) stays the same. So, in the toy dog situation the *alive* property is negated but toy dogs still have four legs, two ears, and a tail, it barks, and so on. These features may need to be appropriately modified (for example the barking is imaginary). However, the uses in the lexicon of techniques from knowledge representation seem open-ended.

6 Acknowledgements

This work was partially supported by the Natural Sciences and Engineering Research Council of Canada grant 0036853, and the Institute for Robotics and Intelligent Systems, a Canadian Network of Centres of Excellence.

References

- [Carnap, 1956] R. Carnap. *Meaning and Necessity*. Chicago University Press, 1956.
- [Cruse, 1986] D. A. Cruse. *Lexical Semantics*. Cambridge University Press, 1986.
- [Etherington *et al.*, to appear] D. W. Etherington, S. Kraus, and D. Perlis. Nonmonotonicity and the scope of reasoning. *Artificial Intelligence*, to appear.
- [Gazdar, 1979] G. J. M. Gazdar. *Pragmatics: Implicature, Presupposition, and Logical Form*. Academic Press, 1979.
- [Grice, 1975] H. P. Grice. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics*, v.3, *Speech Acts*, pages 41–58. Academic Press, 1975.
- [Karttunen and Peters, 1979] L. Karttunen and S. Peters. Conventional implicature. In C-K. Oh and D. A. Dineen, editors, *Syntax and Semantics*, v.11, *Presuppositions*, pages 1–56. Academic Press, 1979.
- [Karttunen, 1973] L. Karttunen. Presuppositions of compound sentences. *Linguistic Inquiry*, 4:169–193, 1973.
- [Karttunen, 1974] L. Karttunen. Presupposition and linguistic context. *Theoretical Linguistics*, 1:181–194, 1974.
- [Mercer and Reiter, 1982] R. E. Mercer and R. Reiter. The representation of presuppositions using defaults. In *Proceedings of the 4th Biennial Conference of the CSCSI/SCEIO*, pages 103–107, 1982.
- [Mercer, 1987] R. E. Mercer. *A Default Logic Approach to the Derivation of Natural Language Presuppositions*. PhD thesis, Dept. of Computer Science, University of British Columbia, 1987.
- [Mercer, 1988a] R. E. Mercer. Solving some persistent presupposition problems. In *Proceedings of the 12th International Conference on Computational Linguistics*, pages 420–425, 1988.
- [Mercer, 1988b] R. E. Mercer. Using default logic to derive natural language presuppositions. In *Proceedings of the 7th Biennial Conference of the CSCSI/SCEIO*, pages 14–21, 1988.
- [Mercer, 1990] R. E. Mercer. Deriving natural language presuppositions from complex conditionals. In *Proceedings of the 8th Biennial Conference of the CSCSI/SCEIO*, pages 114–120, 1990.
- [Reiter, 1980] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [Soames, 1979] S. Soames. A projection problem for speaker presuppositions. *Linguistic Inquiry*, 10:623–666, 1979.
- [Soames, 1982] S. Soames. How presuppositions are inherited: A solution to the projection problem. *Linguistic Inquiry*, 13:483–545, 1982.