Good and Bad News in Formalizing Generalized Implicatures
Author(s): Jacques Wainer and Anthony Maida
Proceedings of the Sixteenth Annual Meeting of the Berkeley Linguistics Society (1990), pp. 530-540

Good and Bad News in Formalizing Generalized Implicatures

Jacques Wainer Anthony Maida
Pennsylvania State University

This paper examines the pros and cons of a model-theoretic formalization of a class of generalized implicatures. We concentrate on generalized implicatures that derive from the assumption that the speaker is observing the maxim of Quantity. There is a family of logics, called nonmonotonic logics, which, at a theoretical level, seem well suited to modeling generalized implicatures. However, when one develops the details, the approach loses its elegance. Difficulties arise because current nonmonotonic logics are not fine-grained enough in the kinds of differential implicatures that they can model. In particular, given two distinct logical forms which are logically equivalent, nonmonotonic logics are forced to derive the same conclusion from each logical form.

Gazdar's (1979) formalization of generalized implicatures has two differences from that examined in this paper. First, Gazdar's approach is centered on the lexical items, whereas our approach is based on model-theoretic concepts. The second difference is that we plan to concentrate on a form of generalized implicature that arises from the assumption that the speaker mentioned all relevant entities in the utterance. This form of implicature does not seem to be easily incorporated into Gazdar's framework.

The next section states why nonmonotonic logic should be an appropriate framework with which to model generalized implicatures. Section 2 discusses a particular form of generalized implicature, the only-implicatures, and we present the similarities between this implicature and circumscription, a particular form of nonmonotonic logic. We then apply the approach to some simple examples with some relative success. Section 3 discusses the major limitations of the approach. We discuss a way of formalizing nonmonotonic logic through semantic concepts which allow us to prove that distinct but logically equivalent logical forms lead to identical conclusions. We then come up with some simple examples for which, we claim, there is no simple way of coming up with distinct non-logically equivalent logical forms. Finally we discuss some of the possible directions for future research in the area of formalizing these kinds of implicatures.

1 The Approach

Our approach is to embed the assumptions stated in Grice's maxim of Quantity into a logic $\mathcal{L}$. We then specify the generalized implicatures of an utterance in terms of logical entailment within $\mathcal{L}$. In other words, the implicatures of an utterance are calculated by performing some logical deduction from the semantic content of the sentence-utterance.

The logic $\mathcal{L}$, if it is to model implicatures, must be nonmonotonic. Monotonicity is the property of a logic that if some formula is entailed from a set of assumptions, then the same formula will also be entailed for any other set of assumptions that contains the original set. If we use the symbol $\models$ to denote logical entailment, Γ to denote the original set of assumptions, and Δ for the additional assumptions, then expression (1) will hold for any monotonic logic.

(1) $\quad$ if $\quad \Gamma \models \alpha \quad$ then $\quad \Gamma \cup \Delta \models \alpha$

Traditional first-order logics, modal logics, intentional logics, and so on, are all monotonic. On the other hand, for a nonmonotonic logic it is possible that

(2) $\quad \Gamma \models \alpha \quad$ but $\quad \Gamma \cup \Delta \not\models \alpha$

Nonmonotonicity is necessary to account for the phenomenon of cancelability of implicatures. For example, the utterance (3.a) carries (3.b) as an implicature. But (3.c) does not carry the same implicature.

(3)
a. Some of the girls are sick.
b. Some of the girls are sick and not all of the girls are sick.
c. Some, and maybe all of the girls are sick.

If we assume that the function [ ] takes as argument a language expression and denotes the expression's semantic content in the language of $\mathcal{L}$, and if the logic $\mathcal{L}$ correctly models implicatures, then the semantic content of (3.b) should be derivable in $\mathcal{L}$ from the semantic content of (3.a). In a more mathematical notation:

(4) $\quad$ { [some of the girls are sick] } $\models_{\mathcal{L}}$ [not all of the girls are sick]

But if one assumes compositionality, then one must acknowledge the semantic content of (3.c) is

(5) $\quad$ [some of the girls are sick] $\wedge$ [maybe all girls are sick]

Thus, nonmonotonicity in a logic is necessary if that logic is to account for the cancelability of implicatures. Furthermore, this framework allows us to model some forms of cancelability derived from context information, as opposed to the form of cancelability by adding extra clauses shown above. For example

(6) $\quad$ The eggs are in the garden or in the living room.

generally implicates that as far as the speaker knows the eggs could be in either place and the speaker does not know which. But given the information that the speaker is the sponsor of an Easter egg hunt, and that he knows where the eggs are, the implicature that he does not know where the eggs are is canceled. Common knowledge is easily represented as formulas in the same logical language one uses to represent the semantic content of sentences. Thus the derivation from the logical form of (6) of some formula that states that the speaker does not know where the eggs are, is canceled with the addition of the assumption that the speaker knows where the eggs are. Other forms of context information, like focus, may also contribute to cancel implicatures. But focus is unlikely to be easily modeled by this approach. This is because it is unclear how to represent focus as a set of formulas in a logical language.

This is the proper place to discuss our assumptions. First we assume that there is such a thing as the literal meaning of utterances. Of this literal meaning, we are interested in the components that correspond to truth functional concepts, which we called above as the propositional content of the utterance. But the most important assumption is that the propositional content, and indeed the literal meaning of an utterance cannot be canceled if the utterance is being used literally. In other words, by adding more clauses to a literal utterance one cannot cancel any part of the original utterance's propositional content. This is the key assumption in the discussion of what is the propositional content of some particular sentences later on.

We will represent the propositional content of an utterance by a formula in some standard logical language, in particular, for most of the examples in this paper we use a first-order language to represent the propositional content of utterances, except for examples where one must be explicit about the state of knowledge of the speaker and then we use a quantified modal logic language.

2 Good News

We can show that the approach suggested in the previous section is indeed able to generate some of the generalized implicatures of utterances. In particular, we concentrate on only-implicatures.

Only-implicatures

We claim that there is a class of generalized implicatures, which we call ONLY-IMPLICATURES, which derives from a consequence of the Quantity maxim, namely that the speaker mentioned all relevant entities in the utterance. Typical examples of only-implicatures are sentences whose implicatures are obtained by adding the particle *only* in some selected places. Let us see some unproblematic examples, where the first sentence is the utterance and the second is what is implicated:

(7) a. John has two children.
b. John has ONLY two children.

(8) a. The flag is white and blue.
b. The flag is ONLY white and blue.

(9) a. Tom and Jerry moved the sofa.
b. ONLY Tom and Jerry moved the sofa.

(10) a. Mary fell from the ferry boat.
b. ONLY Mary fell from the ferry boat.

So for example, in (7.a) the speaker mentions the existence of two of John's children. If one assumes that the speaker mentioned all of the relevant entities, and relevance here seems to mean the entities that are children of John's, then one can conclude the total number of John's children. Actually, all statements above must be qualified by the speaker's knowledge. One should read the explanation above as: since the speaker mentioned the existence of two of John's children, one can assume that that is the total number of John's children as far as the speaker knows. In the remainder of this paper, we will leave out the references to the speaker's knowledge, and the proper qualification must be assumed by the reader. Only when the knowledge of the speaker (or lack of it) becomes important as part of the literal meaning of a sentence, or as part of the implicature, will we explicitly refer to it.

Different aspects of only-implicatures have been discussed in the literature before. For example, the implicature described in (7) above can be explained by the scalar implicature formalizations of Horn (1972) and of Gazdar (1979). Also, only-implicatures seem to be related with Harnish (1976) all-implicatures from where we take example (8) above. But the scope of only-implicatures extends beyond those approaches, as examples (9) and (10) show. In particular, it seems unlikely that Gazdar's formalization, which is centered on lexical items, could be extended to include examples (9) and (10) because there is no lexical item that can be said to be the trigger of the implicature.

However, the existence of only-implicatures as a general phenomenon is debatable (Fauconnier 1990). The idea that all relevant entities have been mentioned would allow for unwanted implicatures, for example, from (7.a) one should conclude that

(11) ONLY John has only two children.

which is clearly not a generalized implicature of (7.a). The extended version of this paper (Wainer and Maida 1990) addresses this point showing that only-implicatures can still be considered as an instance of generalized implicatures that are further canceled by stereotypical knowledge about the number of children one usually has.

The assumption that all relevant objects have been mentioned is also present in the field of formalizing common sense reasoning. In the attempt to provide a formalization to the common sense inference described in (12), McCarthy(1980) created a family of nonmonotonic logics, called CIRCUMSCRIPTION, that capture the assumption that all objects have been mentioned.

(12) Most birds fly.
Penguins do not fly and they are birds.
Tweety is a bird.

Tweety flies

McCarthy's solution is to make explicit a predicate of being abnormal-in-the-aspect-of-flying, which all penguins satisfy. In other words, all birds that are not abnormal-in-the-aspect-of-flying do fly and all penguins are abnormal-in-the-aspect-of-flying. Given that Tweety is a bird, and if we assume that given what is known, abnormal-in-the-aspect-of-flying satisfies the smallest number of birds, then one may conclude that Tweety can fly. More simply, if Tweety is not known to be abnormal-in-the-aspect-of-flying, assume it is not.

Before we discuss the details of circumscription, we would like to point out the connections (at this moment only intuitive) between the problem stated in (12) and Grice's quantity maxim. It seems to be exactly because one assumes that statement 'Tweety is a bird' was as informative as required for the purposes of the problem that one would want to conclude that Tweety flies.

Circumscription

Given a set Γ of assumptions from which one would like to derive conclusions, predicate circumscription is accomplished by adding a second order axiom (called the CIRCUMSCRIPTION SCHEMA) to the set of assumptions 'before' performing the logic derivations. Derivations are then performed in (monotonic) second-order logic. Nonmonotonicity arises from the fact that the schema is added 'after' all assumptions have been collected, and so if the set of assumptions is changed to $\Gamma \wedge \beta$, the circumscription schema itself would be different, and this different axiom could invalidate conclusions that have been derived using the previous schema.

If P is the predicate one wishes to circumscribe, and $\Gamma(P)$ is the set of assumptions, then the circumscription schema will be

(13) $\forall \phi \, (\Gamma(\phi) \rightarrow P \leq \phi)$

where ϕ is a predicate variable of the same arity of P, $\Gamma(\phi)$ is the same set of assumption where the symbol P has been substituted everywhere by ϕ, and $P \leq \phi$ is an abbreviation for

(14) $\forall x \; (P(x) \rightarrow \phi(x))$ where x is a tuple of variables of the appropriate size.

What (13) states is that if the set of conclusions holds for another predicate substituted in the place of P then the new predicate is at least as large as P. Or in other words, it states that P has the smallest extension compatible with Γ.

More complex forms of circumscription have been developed: POINTWISE CIRCUMSCRIPTION (Lifschitz 1986); PARALLEL CIRCUMSCRIPTION and PRIORITIZED CIRCUMSCRIPTION (McCarthy 1980). These formalisms propose different forms in which the minimization of the extension of a predicate is carried out.

Examples

As it was shown above, only-implicatures seem to have strong connections with circumscription. We will now show that circumscription will generate correct only-implicatures for some cases. However, to calculate the implicatures by applying circumscription, one must first have in hand the propositional content of the sentence. It is to the discussion of the propositional content of (15) that we now turn.

(15) John has two children.

As we stated earlier, we will represent the propositional content of sentences by formulas first order logic. We claim that the propositional content of (15) is fully captured by the formula

(16) $\exists xy\ [\ \text{child}(x,\text{John}) \land \text{child}(y,\text{John}) \land x \neq y\ ]$

which states that there are two different entities, each of which is a child of John. What should be made clear is that the formula does not specify the exact number of John's children, but it sets a lower bound to the number of John's children as two. Or in other words, the propositional content of (15) is a formula that one would associate with the import of 'John has at least two children.' To defend this claim we must show that there is no stronger formula that can capture the propositional content of (16).

One could reasonably hypothesize that the propositional content of (15) is that John has exactly two children, represented by expression (17) below.

(17) $\exists xy\ [\ \text{child}(x,\text{John}) \land \text{child}(y,\text{John}) \land x \neq y \land \forall z\ [\ \text{child}(z,\text{John}) \rightarrow (x = z \lor y = z)\ ]\ ]$

The reason why (17) seems too strong to be the propositional content of (15) derives from a compositionality argument, coupled with the assumption that literal meaning is not cancellable. In particular, the literal meaning of the first clause of expression (18) below

(18) John has two children, in fact he has three.

should be identical to the literal meaning of (15). The propositional content of the second clause of (18) above should be analogous to (17) except that the expression should capture the fact that there are three, instead of two children. If we assume that the truth functional import of *in fact* is the same as the logical *and* then the propositional content of (18) should be the same as that of expression (19), but (19) is inconsistent and (18) is not.

(19) John has exactly two children, in fact he has exactly three.

The claim that *in fact* has the same truth functional import as *and* is debatable (Wilensky 1990). The expression *in fact* in some way asserts both clauses it links, although it has some other pragmatic imports. But since literal meaning cannot be canceled, none of these extra factors will cancel the fact that both clauses are being asserted, or in other words, *in fact* has at least the same truth functional import as *and*. Actually, the main source of limitations for this assumption is that *in fact* is being modeled by a symmetric operator when it is really not a symmetric connector. The consequence is that the theory will make incorrect predictions for anomalous expressions like

(20) *John has three children, in fact he has two.

We have shown that the propositional content of (15) cannot be (17), and that it must be (16). This result is not as counter intuitive as it may seem in a first analysis. The fact that (17) is not the propositional content of (15) states that the fact that John has only two children is really an implicature and not an entailment of (15).

Given that we accept that (16) is the propositional content of (15), we must now specify a logic such that (15) would entail the proper only-implicature. This is accomplished by circumscribing the predicate $\lambda(x)$child(x, John) (being a child of John's) in (16). The formula

(21) $\exists xy$ [child(x,John) $\wedge$ child(y,John) $\wedge$ x $\neq$ y
$\wedge$ $\forall z$ [child(z,John) $\rightarrow$ (x = z $\vee$ y = z)]]

which roughly states that John has only two children, will be derivable from the circumscription schema.

To be more precise, what is done is the circumscription of a dummy predicate symbol Q defined to be equivalent to $\lambda(x)$child(x, John). In this case the set of assumptions is

(22) $\forall z$ [Q(z) $\leftrightarrow$ child(z,John)] $\wedge$
$\exists xy$ [child(x,John) $\wedge$ child(y,John) $\wedge$ x $\neq$ y]

and Q is circumscribed in (22) allowing the extension of child to vary as needed. From now on, when we mention the circumscription of a complex predicate, one that is constructed using the λ-operator, we mean a circumscription carried along the lines described above.

Implicature derived from circumscription are nonmonotonic, as the example below will show. The propositional content of

(23) John has two children, in fact he has three.

will be equivalent to

(24) $\exists xy$ [child(x,John) $\wedge$ child(y,John) $\wedge$ x $\neq$ y] $\wedge$
$\exists xyz$ [child(x,John) $\wedge$ child(y,John) $\wedge$ child(z,John)
$\wedge$ x $\neq$ y $\wedge$ x $\neq$ z $\wedge$ y $\neq$ z]

Again, circumscribing $\lambda(x)$child(x, John) in (24) above would result in a formula that states John has only three children, which is the correct implicature.

A similar procedure when applied to examples below

(25) a. The flag is white and blue.
b. John moved the sofa and the bed.

would generate the correct implicatures in (26)

(26) a. The flag is only white and blue.
b. John moved only the sofa and the bed.

In the first example what is circumscribed is $\lambda(\mathtt{x})\mathtt{color}(\mathtt{x},\mathtt{y})$ (being the color of something). In the second example the predicate $\lambda(\mathtt{x})\mathtt{move}(\mathtt{John},\mathtt{x})$ (being moved by John) is circumscribed.

On the other hand, the limitations of those results are clear. The ad hoc way in which the predicate to be circumscribed has to be picked must be explained and further research is being done in the direction of automatizing this choice. And the results are very sensitive to that choice. For example, going back to (15) and its logical form (16), if instead of circumscribing $\lambda(\mathtt{x})\mathtt{child}(\mathtt{x},\mathtt{John})$ we had circumscribed the predicate symbol `child` the resulting conclusion would have been a formula that has the same import of (27) below.

(27) There is only one parent, which is John, and he has exactly two children.

3 Bad News

Although the examples developed in the previous section suggest that the approach has some potential, we have found examples where the idea of modeling implicatures as entailment within a nonmonotonic logic cannot succeed, for any logic of interest. The reason for the failure is that circumscription (actually a larger class of nonmonotonic logics) is insensitive to logical equivalence. In other words, two different formulas that are logically equivalent will entail the same set of conclusions no matter what predicate is being circumscribed or what version of circumscription is being used.

To show that we must first explain this larger class on nonmonotonic formalisms, which contains all forms of circumscription, and then prove that any logic in this class will be insensitive to logical equivalence.

All forms of circumscription can be classified as semantically based logics. This is to say that besides the axiomatic definition one can also give a model-theoretic characterization of the various forms of circumscription. Those model-theoretic characterizations are all based in an idea, first developed by Shoham (1987), of defining an ordering relation among the class of models for a set of formulas. This ordering relation, called PREFERENCE, defines a partial order of models, and allows one to define MAXIMALLY PREFERRED MODELS as models for which there is no other model that is 'more preferred'.

LOGICAL ENTAILMENT is then defined as: a set of formulas Γ logically entails β if and only if all maximally preferred models for Γ also satisfy β. The traditional definition of logical entailment would require that all models that satisfy Γ must also satisfy β. Different logics are created for different definitions of the preference relation $\Im$. For example, the preference relation for the predicate circumscription of a predicate symbol P, would state that M_1 is preferred to M_2 if they agree with each other in the denotation of every constant symbol and every predicate symbol, except that the extension of P in M_1 is smaller than in M_2.

Given the above description of the semantically-based nonmonotonic logics, which includes all forms of circumscription, we can state the theorem that logically equivalent formulas would yield the same conclusion under any of those logics.

The proof is based on the fact that logically equivalent formulas are satisfied by the same models. Therefore, under any preference criteria, the set of maximally preferred models are the same for both formulas. Since the logical entailments of a formula are exactly the formulas that are satisfied by all maximally preferred models, and logically equivalent formulas have the same set of maximally preferred models, then both formulas will have the same set of entailments.

Given these facts we can now discuss some examples where the sentences would have equivalent logical forms but have different implicatures. Let us compare the sentences (28) and their implicatures in (29)

(28) a. John has two children.
 b. John has at least two children.

(29) a. John has only two children.
 b. John has two children and possibly more.

We have argued that the propositional content of (28.a) is captured by expression (16), which also seems to be the propositional content of (28.b). In fact, (16) captures exactly what one would expect to the the content of (28.b): a statement on the lower bound of the number of John's children. But if this is so, then both sentences in (28) have logically equivalent literal meanings and there is no nonmonotonic logic that can derive the correct implicature for both sentences simultaneously.

It is possible that we have given up on expression (28.b) too easily and that its logical form does contain more information than that shown in (16). We have not accounted for the contribution of the phrase *at least* in the literal meaning of the sentence. Perhaps, it means that the speaker considers possible that John has more than two children, and this should be part of the literal meaning of the sentence. In other words, the propositional content of (28.b) would be captured by the import of

(30) The speaker knows that there are (at least) two things that are John's children and it is possible considering his knowledge that there is (at least) a third child.

The extended version of this paper (Wainer and Maida 1990) shows that if (30) is represented by a modal logic formula (Hintikka 1962), then

(31) John has at least two children, in fact he has exactly two.

would be contradictory.

Thus, under reasonable theories of propositional content the pair of sentences in (28) would yield logically equivalent logical forms. Thus it is impossible to construct a nonmonotonic logic that simultaneously derive the correct implicature for each sentence..

Other sentences which under a reasonable theory of propositional content would have semantic interpretations which are logically equivalent but whose correspondent sentence-utterances carries different implicatures is the pair below:

(32) a. John has two or three children.
b. John has two or more children.

In a nutshell, the problem with the sentences above is that *three* is translated to 'there are (at least) three things that satisfies ...' and *more*, which we take to be an ellipsis of 'more than two', would have an equivalent translation.

4 Conclusions

We presented an approach to modeling generalized quantity implicatures as entailment in some nonmonotonic logic. This approach seems natural and has potential. We have shown that circumscription does derive some of the correct only-implicatures. Some examples of clausal implicatures, in particular the interaction between modal operators like *knowledge* and logical connectives, although not discussed in this paper are also being investigated by the authors.

On the other hand, we have identified a major limitation of the approach: its inability to distinguish between logically equivalent formulas. We have pointed out examples where a reasonable interpretation of the semantic content of some sentences generates logically equivalent logical forms for sentences that carry different implicatures.

Future research points in the direction of incorporating the calculation of generalized implicatures in the semantic interpretation itself along the lines of Karttunen and Peters (1979) work in conventional implicatures. There is evidence that the derivation of generalized implicatures is centered more on lexical and syntactical items than on semantic concepts. For example, the presence of the expression *at least* in (28.b) is responsible for a different set of implicatures in comparison with (28.a). Other evidence comes from the examples below:

(33) a. John sees two children.
b. John seeks two unicorns.

The implicature for both cases is very similar:

(34) a. John sees only two children.
b. John seeks only two unicorns.

but the logical forms for sentences (33) are much different. It is more reasonable to assume that the implicature derivation process is in the semantic interpretation function. This is because the presence of the word *two*, independent on whether it appears within an intentional context or not, generates the implicature. This theory seems more plausible than to assume that some logic, working from very different logical forms, would derive such similar implicatures.

References

Fauconnier, Gilles. 1990. Personal communication, 16th Annual Meeting of the BLS.

Gazdar, Gerald. 1979. *Pragmatics: Implicature, Presupposition and Logical Form.* Academic Press, New York.

Grice, Paul. 1975. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics 3: Speech Acts*, Academic Press, New York.

Harnish, Robert M. 1976. Logical form and implicature. In T. Bever, J. Katz, and T. Langendoen, editors, *An Integrated Theory of Linguistic Ability*, pages 313–392, Thomas Y. Crowell Company, New York.

Hintikka, Jaakko. 1962. *Knowledge and Belief.* Cornell University Press.

Horn, L. 1972. On the semantic properties of the logical operators in English. Unpublished doctoral dissertation, UCLA, Los Angeles.

Karttunen, Lauri and Peters, Stanley. 1979. Conventional implicature. In Choon-Kyu Oh and David A. Dinneen, editors, *Syntax and Semantics: Presupposition*, Academic Press, New York.

Lifschitz, V. 1986. Pointwise circumscription: preliminary report. In *Proc. American Assoc. for Artificial Intelligence*, pages 406–410, Philadelphia, PA.

McCarthy, John. 1980. Circumscription—a form of nonmonotonic reasoning. *Artificial Intelligence*, (13):27–39.

Shoham, Y. 1987. A semantical approach to nonmonotonic logics. In *Proc. of the Tenth International Joint Conference on Artificial Intelligence*, pages 388–392.

Wainer, Jacques and Maida, Anthony. 1990. Good and Bad News in Formalizing Generalized Implicatures (extended version). Technical Report CS–90–11 , The Pennsylvania State University, Department of Computer Science, 333 Whitmore Lab, University Park, PA 16802.

Wilensky, Robert. 1990. Personal communication, 16th Annual Meeting of the BLS.