Local Pragmatics

Jerry R. Hobbs and Paul Martin
Artificial Intelligence Center
SRI International

Abstract

The outline of a unified theory of local pragmatics phenomena is presented, including an approach to the problems of reference resolution, metonymy, and interpreting nominal compounds, and the TACITUS system embodying this theory is described. The theory and system are based on the use of a theorem prover to draw the appropriate inferences from a large knowledge base of commonsense and technical knowledge. Issues of control and minimality are discussed.

1 The Problems

In the messages about breakdowns in machinery that are being processed by the TACITUS system at SRI International, we find the following sentence:

> We disengaged the compressor after the lube oil alarm.

This sentence, like virtually every sentence in natural language discourse, confronts us with difficult problems of interpretation. First, there are the reference problems; what do "the compressor" and "the lube oil alarm" refer to. Then there is the problem of interpreting the implicit relation between the two nouns "lube oil" (considered as a multiword) and "alarm" in the nominal compound "lube oil alarm". There is also a metonymy that needs to be expanded. An alarm is a physical object, but "after" requires events for its arguments. We need to coerce "the lube oil alarm" into "the sounding of the lube oil alarm". There is the syntactic ambiguity problem of whether to attach the prepositional phrase "after the lube oil alarm" to "the compressor" or to "disengaged". All of these problems we have come to call problems in "local pragmatics". They seem to be specifically linguistic problems, but the traditional linguistic methods in syntax and semantics have not yielded solutions of any generality.

The difficulty, as is well-known, is that to solve these problems we need to use a great deal of arbitrarily detailed general commonsense and domain-specific technical knowledge. A theory of local pragmatics phenomena must therefore be r. theory about how knowledge is used. The aim of our research has been to develop a unified theory of local pragmatics, encompassing reference resolution, metonymy, the interpretation of nominal compounds and other implicit and vague predicates, and the resolution of syntactic, lexical, and quantifier scope ambiguities, based on the drawing of appropriate inferences from a large knowledge base.

The TACITUS system is intended to embody that theory. Its specific aim is the interpretation of casualty reports (casreps), which are messages in free-flowing text about breakdowns in mechanical devices. More broadly, however, our aim is to develop general procedures, together with the underlying theory, for using commonsense and technical knowledge in the interpretation of written (and spoken) discourse regardless of domain.

The TACITUS system has four principal components. First, a syntactic front-end, the DIALOGIC system, translates sentences of a text into a logical form in first-order predicate calculus, described in Section 3. Second, we are building a knowledge base, specifying large portions of potentially relevant knowledge encoded as predicate calculus axioms (Hobbs et al., 1986). Third, the TACITUS system makes use of the KADS theorem prover, developed by Mark Stickel (Stickel, 1982). Finally, there is a pragmatics module, including a local pragmatics component, which uses the theorem prover to draw appropriate inferences from the knowledge base, thereby constructing an interpretation of the text.

To interpret a sentence, the local pragmatics component determines from the logical form of the sentence what interpretation problems need to be solved. Logical expressions are constructed for each local pragmatics phenomenon the sentence exhibits, and the proofs of these expressions constitute the interpretation of the sentence. Where there is more than one interpretation, it is because there is more than one proof for the expressions. Interpretation is thus viewed as crucially involving deduction.

In Section 2, we describe the three of the phenomena we are addressing—reference, metonymy, and nominal compounds. For each, we describe the expression that needs to be proved, and for the last two, we describe how current standard techniques can be seen as special cases of our general approach.

2 Local Pragmatics Phenomena

2.1 Reference

Entities are referred to in discourse in many guises—proper nouns, definite, indefinite, and bare noun phrases of varying specificity, pronouns, and omitted or implicit arguments. Moreover, verbs, adverbs, and adjectives can refer to events, conditions, or situations. The problem in all of these cases is to determine what is being referred to. Here we confine ourselves to definite noun phrases, although TACITUS handles the other cases in a manner described briefly in Section 4.

In the sentence

The alarm sounded.

the noun phrase "the alarm" is definite, and the hearer is therefore expected to be able to identify a unique entity that the speaker intends to refer to. Restating this in theorem-proving terminology, the natural language system should be able to prove constructively the expression

$$(\exists x)alarm(x)$$

It must find an *x* which is an alarm in the domain model. If it succeeds, it has solved the reference problem.[1]

2.2 Metonymy

Metonymy is very common in discourse; few sentences lack examples. Certain functions frequently provide the required coercions. Wholes are used for parts, tokens for types, and objects for their names. Nunberg (1978), however, has shown that there is no finite set of possible coercion functions. The relation between the explicit and implicit referents can be virtually anything.

Every morpheme in a sentence conveys information that corresponds to a predication, and every predicate imposes *selectional constraints* on its arguments. Since entities in the text are generally the arguments of more than one predicate, there could well be inconsistent constraints imposed on them. To eliminate this inconsistency, we interpose, as a matter of course, another entity and another relation between any two predications. Thus, when we encounter in the logical form of a sentence

$$\ldots \wedge after(e_0, a) \wedge alarm(a)$$

we assume that what is intended is really

$$\ldots \wedge after(e_0, z) \wedge q(z, a) \wedge alarm(a)$$

for some entity *z* and some relation *q*. The predication *q(z, a)* functions as a kind of buffer, or impedence match, between the explicit predications with their possibly inconsistent constraints. In many cases, of course, there is no inconsistency. The argument satisfies the selectional constraints imposed by the predicate. In this case, *z* is *a* and *q* is identity. This is the first possibility tried in the implemented system. Where this fails, however, the problem is to find what *z* and *q* refer to, subject to the constraint, imposed by the predicate *after,* that *z* is an event.

TACITUS thus modifies the logical form to

$$\ldots \wedge after(e_0, z) \wedge q(z, a) \wedge alarm(a)$$

and for an interpretation, the expression that must be proved constructively is

$$(\exists z, q, a)event(z) \wedge q(z, a) \wedge alarm(a)$$

Find an event *z* bearing some relation *q* to the alarm.

The most common current method for dealing with metonymy, e.g., in the TEAM system (Grosz et al., 1985), is to specify a small set of possible coercion functions, such as *name-of.* This method can be captured in the present framework by treating *q* not as a predicate variable, but as a predicate constant, and expressing the possible coercions in axioms like the following:

$$(\forall x, y)name(x, y) \supset q(x, y)$$

That is, if *x* is the name of *y,* then *y* can be coerced to *x.* This is the method we have implemented in our initial version of the TACITUS system.

2.3 Nominal Compounds

To interpret a nominal compound, like "lube oil alarm", it is necessay to discover the implicit relation between the nouns.[2] Some relations occur quite frequently in nominal compounds—*part-of*, *location, purpose.* Moreover, when the head noun is relational, the modifier noun is often one of the arguments of the relation. Levi (1978) argued that these two cases encompassed virtually all nominal compounds. However, Downing (1977) and others have shown that virtually any relation can occur. A lube oil alarm, for example, is an alarm that sounds when the pressure of the lube oil drops too low.

To discover the implicit relation, one must prove constructively from the knowledge base the existence of some possible relation, which we may call *nn,* between the entities referred to by the nouns:

$$(\exists x, y, nn)alarm(x) \wedge lube\text{-}oil(y) \wedge nn(y, x)$$

In our framework, we can implement the approach that hypothesiszes a small set of possible relations, by taking *nn* to be not a predicate variable but a predicate constant, and encoding the possibilities in axioms like

[1] In this paper we ignore the problem of the uniqueness of the entity referred to. A hint of our approach is this: If the search for a proof is heuristically ordered by salience, then the entity found will be the uniquely most salient.

[2] Some nominal compounds can of course be treated as single lexical items. This case is not interesting and is not considered here.

$$(\forall x, y) part(x, y) \supset nn(y, x)$$

For example, if a blade x is a part of a fan y, then "fan blade" is a possible nominal compound.

To deal with relational nouns, such as "oil sample", we encode axioms like

$$(\forall x, y) sample(x, y) \supset nn(y, x)$$

This tells us that if x is a sample of oil y, then x can be referred to by the nominal compound "oil sample".

Finin (1980) argues that one of the most common kinds of relations is one that involves the function of the referent of the head noun. The function of a pump is to pump a fluid, so "oil pump" is a possible nominal compound. This can be encoded in axioms of the pattern

$$(\forall x, y, e) function(e, x) \wedge p'(e, x, y) \supset nn(y, x)$$

That is, if e is the function of x where e is the situation of x doing something p to y, then there is a possible nn relation between y and x.

As with metonymy, in our initial version of TACITUS, it was the standard, restricted method that we implemented. This is because we wanted to make sure we were not losing ground in seeking a general solution. Nevertheless, our approach allows us to begin experimenting with the general solution to the nominal compound problem, where the implicit relation can be anything at all.

3 Order of Interpretation

DIALOGIC produces a logical form for the sentence in something like a first-order logic but encodes grammatical subordination relations as well as predicate-argument relations. It is "ontologically promiscuous" in that events and conditions are reified. A slightly simplified version of the logical form for sentence (1) is (2):

(1) The lube oil alarm sounded.

(2) $past([e_1 \mid sound'(e_1, [a_1 \mid alarm(a_1) \wedge nn([o_1 \mid lube\text{-}oil(o_1)], a_1)])])$

"|" can be read "such that" or "where", so that a paraphrase of this formula would be "In the past there was an event e_1 which was a sounding event by a_1 where a_1 is an alarm and there is an nn relation between a_1 and o_1 such that o_1 is lube oil."

In general, the logical form of a sentence is a "proposition". A proposition is a predicate applied to one or more arguments. An argument is either a variable or a "complex term". A complex term is a variable, followed by a "such that" sign, followed by a "restriction". (Complex terms are surrounded by square brackets for readability.) A restriction is a conjunction of propositions.

This notation can be translated into a notation using four-part quantifier structures by successively applying the following transformation:

$$p([x \mid q(x)]) \Rightarrow (\exists x\ q(x)\ p(x))^3$$

Since interpretation involves solving a number of problems, or proving a number of expressions, a question is raised. In which order should we try to solve them? A naive answer would be to try to solve them "from the inside out". Before trying to find the lube oil *alarm*, we should try to find the lube oil the alarm is an alarm *for*. Before checking that the lube oil alarm obeys the selectional constraints imposed by "sound", we should learn as much as we can about the lube oil alarm; in particular, what "the lube oil alarm" refers to.

This means that given the logical form (2), we should solve the local pragmatics problems in the following order:

1. Find the reference of o_1, the lube oil. Prove
$$(\exists o_1) lube\text{-}oil(o_1)$$
2. Given that, find the implicit relation nn encoded in the nominal compound. If o_1 was resolved to O, then prove
$$(\exists nn) nn(a_1, O)$$
3. Given that, find the reference of a_1, the alarm. If nn is the relation for, then prove
$$(\exists a_1) alarm(a_1) \wedge for(a_1, O)$$
4. Given that, check the predicate-argument congruence of *sound* applied to a_1. If a_1 was resolved to A and *sound* requires its argument to be a physical object, then prove
$$(\exists q, z) physical\text{-}object(z) \wedge q(z, A)$$

Unfortunately, this order will not always work. Information relevant to the solution of any of these local pragmatics problems can come from the solutions of any of the others. This can be seen particularly in problems 2 and 3. Sometimes we will need to have resolved the reference to have enough information to solve the nominal compound problem, and sometimes conversely.

Thus, in a more sophisticated approach, we would construct a single expression to be proved, encoding what is required for *all* of the local pragmatics problems. For sentence (1), the expression would be

$$(\exists x, q, a_1, nn, o_1) physical\text{-}object(x) \wedge q(x, a_1) \wedge alarm(a_1) \wedge nn(a_1, o_1) \wedge lube\text{-}oil(o_1)$$

The conjuncts of this expression could be proved in any order. The inside-out order is only one possibility; it can be taken as an indication of how much work the theorem-prover should devote to proofs of the various conjuncts, and how early. The difficulty with this approach, however, is that it is difficult to get partial results in cases of failure.

We are currently using a comprimise between these two orders—a fail-soft, inside-out order. As we proceed inside

[3] Quantifiers other than existentials are ignored in this paper. For the treatment we intend to give them, see Hobbs (1983).

out, at each step the theorem-prover is given the full expression built up to that point. However, the expression has as an antecedent the instantiations of what was proven in earlier steps. Thus, in step 4 in the example, the expression is

$$lube\text{-}oil(O) \wedge alarm(A) \wedge for(A,O) \supset$$
$$(\exists z, q, a_1, nn, o_1) physical\text{-}object(z) \wedge q(z, a_1)$$
$$\wedge alarm(a_1) \wedge nn(a_1, o_1) \wedge lube\text{-}oil(o_1)$$

Those prior instantiations consistent with higher constraints will be proven immediately from the antecedent, and new proofs will need to be discovered only for those which are inconsistent.[4]

4 Other Issues

Syntactic Ambiguity: Within this framework we have also implemented a treatment of the most common syntactic ambiguities. Attachment ambiguity problems are translated into constrained reference problems by introducing an existentially quantified variable representing, say, the subject of the prepositional phrase, constrained by a disjunction listing the possible attachment sites.

Given and New: A problem that arises in the casreps in a severe form, since determiners are often omitted, is the problem of determining whether an entity referred to is definite or indefinite, that is, given or new. In

Metal particles in oil sample and strainer.

"oil sample" is indefinite (or new) while "oil strainer" is definite (or given). This problem occurs in general discourse as well, especially for nonnominal reference. Our approach is to assert the existence of each entity, but with a cost determined by the type of reference—in increasing order, indefinite, nonnominal, bare, or definite. If this assertion (called a *referential implicature)* is required in the proof that constitutes the interpretation, the entity is new; if not, it is given.

Minimality: Axioms can be assigned a cost, depending upon their salience. High salience, low cost axioms would then be tried first. Short proofs are naturally tried before long proofs. Thus, a cost depending on salience and length is associated with each proof, and hence with each interpretation. Where, as usually happens, there is more than one possible interpretation, the better interpretations are supported by less expensive proofs. A second criterion for good interpretations is to favor the minimal solution in the sense that the fewest new entities and relations needed to be hypothesized. For example, the argument-relation pattern in nominal compounds, as in "oil sample", is minimal in that no new implicit relation need be hypothesized; the one already given by the head noun will do. In metonymy, the identity coercion is favored for the same reason, and shorter coercions are favored over longer ones. These ideas at least give us a start on the very difficult problem of choosing the best interpretation.

[4] This technique is due to Mark Stickel.

Acknowledgements

The authors have profited from discussions with Mark Stickel, Doug Edwards, Bill Croft, Fernando Pereira, Ray Perrault, Stu Shieber, and Mabry Tyson about this work. The research was funded by the Defense Advanced Research Projects Agency under Office of Naval Research contract N00014-85-C-0013, as part of the Strategic Computing Program.

References

[1] Downing, Pamela, 1977. "On the Creation and Use of English Compound Nouns", *Language,* vol. 53, no. 4, pp. 810-842.

[2] Finin, Timothy, 1980. "The Semantic Interpretation of Nominal Compounds", Report T-96, Coordinated Science Laboratory, University of Illinois, Urbana, Illinois, June 1980.

[3] Crosz, Barbara J., Douglas E. Appelt, Paul Martin, Fernando C. N. Pereira and Lorna Shinkle, 1985. "The TEAM Natural-Language Interface System", Final Report, Project 4865, Artificial Intelligence Center, SRI International, Menlo Park, California.

[4] Hobbs, Jerry R., 1983. "An Improper Treatment of Quantification in Ordinary English", *Proceedings, 21st Annual Meeting of the Association for Computational Linguistics,* Cambridge, Massachusetts, pp. 57-63.

[5] Hobbs, Jerry R., William Croft, Todd Davies, Douglas Edwards, and Kenneth Laws, 1986. "Commonsense Metaphysics and Lexical Semantics", *Proceedings, 24th Annual Meeting of the Association for Computational Linguistics,* New York, June 1986., pp. 231-240.

[6] Levi, Judith, 1978. *The Syntax and Semantics of Complex Nominals,* Academic Press, New York.

[7] Nunberg, Geoffery, 1978. "The Pragmatics of Reference", Ph. D. thesis, City University of New York, New York.

[8] Stickel, Mark E., 1982. "A Nonclausal Connection-Graph Theorem-Proving Program", *Proceedings, AAAI-82 National Conference on Artificial Intelligence,* Pittsburgh, Pennsylvania, pp. 229-233.