Capturing the deep meaning of texts through deduction and inference

by F. Antonacci
C. M. Calamani

One of the main problems in the computer analysis of natural language is understanding sentences beyond a surface level, i.e., making inferences about likely circumstances and drawing plausible conclusions. At the first level, a natural-language-understanding system can answer simple and trivial questions; in order to extend the domain of possible questions that it can answer, the system must make presuppositions and recognize implications that depend on certain events (also called actions). The IBM Rome Scientific Center has developed a prototype system that is able to make inferences about what might be true. This system has been integrated with a text-understanding system (System N), also developed at Rome.

Introduction

We believe that the ultimate goal of a text-understanding system is to produce a "deep" representation, but the methods by which this representation should be derived are unclear and not generally accepted at the present state of the art. Deep knowledge representation intrinsically lacks generality; depending upon the specific focus of investigation, different and highly context-dependent knowledge aspects can be highlighted. An understanding of deep structure is crucial in capturing all possible aspects of meaning in a sentence. Starting with generative semantics (now considered unsuccessful), the problem has been addressed in several ways. Among these, two approaches have emerged in artificial intelligence (AI) for representing real-world knowledge: *frames* and *scripts.* Minsky's frames [1] encode elements of knowledge needed in a common-sense reasoning system, in a structured and flexible way. However, this theory has not been completely developed and has not obtained satisfactory results, principally because the capability of inserting frames into other frames, and having each of them look like a set of expectations, creates problems and conflicts. The frame theory is probably, at present, better developed by Schank's group under the name of *script* [2]. The real problem involves the heavy hand-coding of the common knowledge in a "script" representation.

In other approaches [3], the phases of superficial and deep comprehension of a text are not separated, but problems arise when the relevant information is connected not to just one or more words in the text but rather to the meaning of the whole sentence.

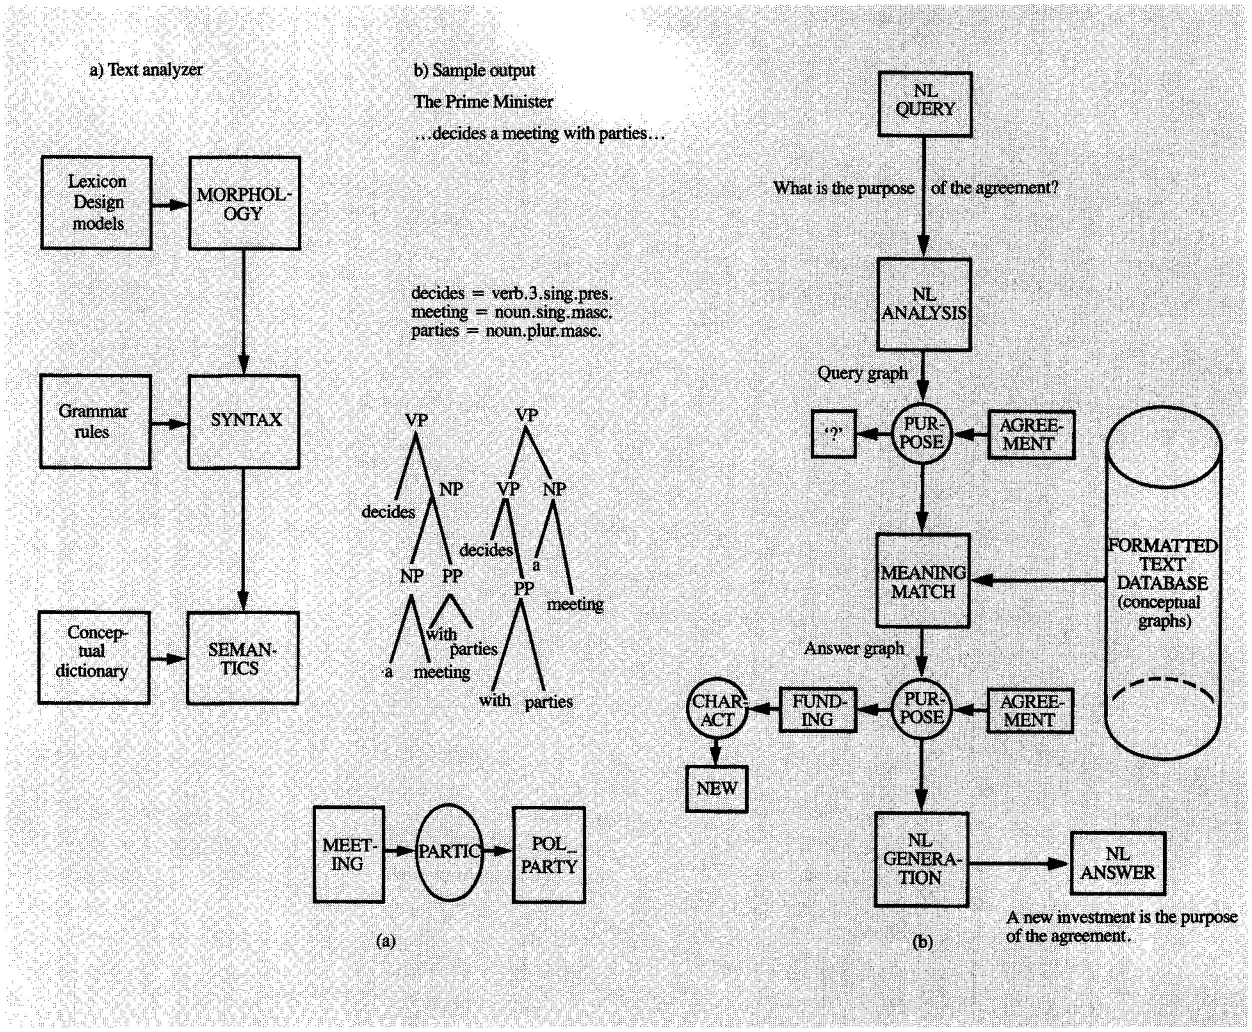

Figure 1

Scheme of the text analyzer and query answering system: (a) text analysis; (b) query processing.

At the IBM Rome Scientific Center, a prototype able to derive common knowledge by means of logical inferences [4] has been developed. The process works on the semantic representation of the texts produced by a system of natural-language understanding, System N [5]. System N analyzes short texts (e.g., press agency news on finance and economics), stores their meaning representation in a knowledge base (KB), and permits the KB to be queried and generate answers in Italian.

The semantic phase of System N is based on a surface semantics, which provides a general and clear encoding of natural-language utterances; in any case, surface graphs are a first step toward the derivation of a "deeper" representation. The inference process embedded in System N allows questions on the deep meaning of the analyzed texts to be formulated by means of inference rules which are defined in a simple way.

The following sections present some of the features of the prototype. The first is a brief overview of the basic components of System N. The remaining sections describe the inferential capability of the prototype, showing by examples how the inference rules are applied to deduce deeper knowledge. The last part of this paper describes our approach to the problem of uncertainty in our text-understanding system.

Overview of System N

System N is a sentence analyzer in the sense that it analyzes complex sentences from a database of press agency releases and gives answers in Italian to a wide spectrum of NL (natural-language)-type questions about the analyzed sentences [6]. In this sense, it can be considered an important step toward an intelligent information retrieval system [7]. **Figure 1** shows the system architecture.

- *Analysis*

The understanding of texts is achieved through three phases of analysis: morphology, syntax, and semantics.

The *morphology* [8, 9] phase analyzes each word of the text, retrieving its morphologic features (*gender, number, tense*, etc.), its syntactic category (*noun, adverb*, etc.), and the lemma from which the word is derived.

The primary components of the morphologic analyzer are 1) a context-free (CF) grammar, to describe the rules of word derivation; 2) a lexicon describing word components, implemented by logic formulae; and 3) other specific CF grammars, to perform the morphosyntactic analysis, recognizing fixed and variable sequences of words, such as idioms, data and number expressions, comparative and superlative forms of adjectives, and compound tenses of verbs. At the current state of implementation, the system has a lexicon of 8000 *elementary lemmata* (root forms without prefixes and suffixes), and allows for full coverage (100%) of the analyzed domain.

The *syntax* [10, 11] phase states the syntactic dependencies among the different words in the sentence and gives one or more representations of the sentence structure through syntactic trees.

The important aspects of the syntactic analyzer are 1) the use of an *attribute grammar* to specify rules on word sequencing and forms agreement, and 2) the use of *look-ahead sets* and early *semantic tests* to avoid the combinatorial explosion of parsing trees. Currently, the grammar has about 150 production rules. A meta-analyzer called MEGA (MEta-analyzer for Attribute Grammars) was developed to write production rules for attribute grammars and to generate automatically and efficiently the corresponding top-down, recursively descendent parser. Estimated coverage of the corpus is 80%, obtained by parsing 1000 press releases randomly selected from the corpus.

The *semantics* [12] phase represents the final step and the most complex part of a natural-language-processing (NLP) system. It resolves syntactic ambiguities and recognizes semantic relationships among words, producing a representation of the sentence meaning through the formalism of conceptual graphs [13].

According to this model, a sentence is represented by a graph of concepts and conceptual relations. To generate a conceptual graph, the semantic interpreter uses

1. A pattern/interpretation table, called SS (syntax to semantics) rules, associating with each syntactic pattern a possible semantic interpretation. A semantic interpretation is the name of a *conceptual relation* that could express the nature of the semantic relation between two words, or phrasal patterns.
2. A semantic lexicon containing for each *word sense* a detailed list of *use types*, called *surface semantic patterns* (SSPs). About 1300 concepts were defined manually by looking at word occurrences in contexts. This semantic lexicon has been enriched with other concepts by using a conceptual semiautomatic acquisition module [14].
3. A hierarchy with various levels of abstraction, where the concepts have been classified; it allows the system to generalize the graph representation of a word.

- *Question answering*

The question-answering module comprises

- A *question analyzer*, which is the same one used for declarative sentences except that during the semantic analysis the node(s) of the conceptual graph corresponding to question pronouns are replaced by temporary "dummy" nodes.
- An *answer retriever*, matching the question graph with the knowledge base and selecting the matched conceptual graph.
- An *answer generator*, producing an answer in Italian that is based on the selected conceptual graph.

Some considerations regarding the knowledge representation model

According to Sowa's model [13], a conceptual graph represents information concerning events, but nothing about the consequences of those events. This means that each graph reflects only the surface semantics of a text, not the deeper one. Let us consider, for example, the sentence represented by the graph in **Figure 2**:

Nel 1987 il gruppo Eni ha venduto alla Marzotto la Lanerossi per 168 miliardi.
[*In 1987 Eni group sold Lanerossi to Marzotto for 168 thousand million lire.*]

From this representation, it is deducible (but not explicit) that Lanerossi belongs to Marzotto, that Marzotto bought Lanerossi from gruppo Eni in 1987, and that Lanerossi has belonged to Marzotto for five years.

The lack of implicit information in the semantic representation of a text becomes critical above all in a question-answering phase: System N can be questioned only about information belonging to the surface semantics of analyzed texts.

Looking at the displayed sentence, we can ask, for example, *Who did sell Lanerossi? When did Eni group sell Lanerossi? To whom did Eni group sell Lanerossi?* But the

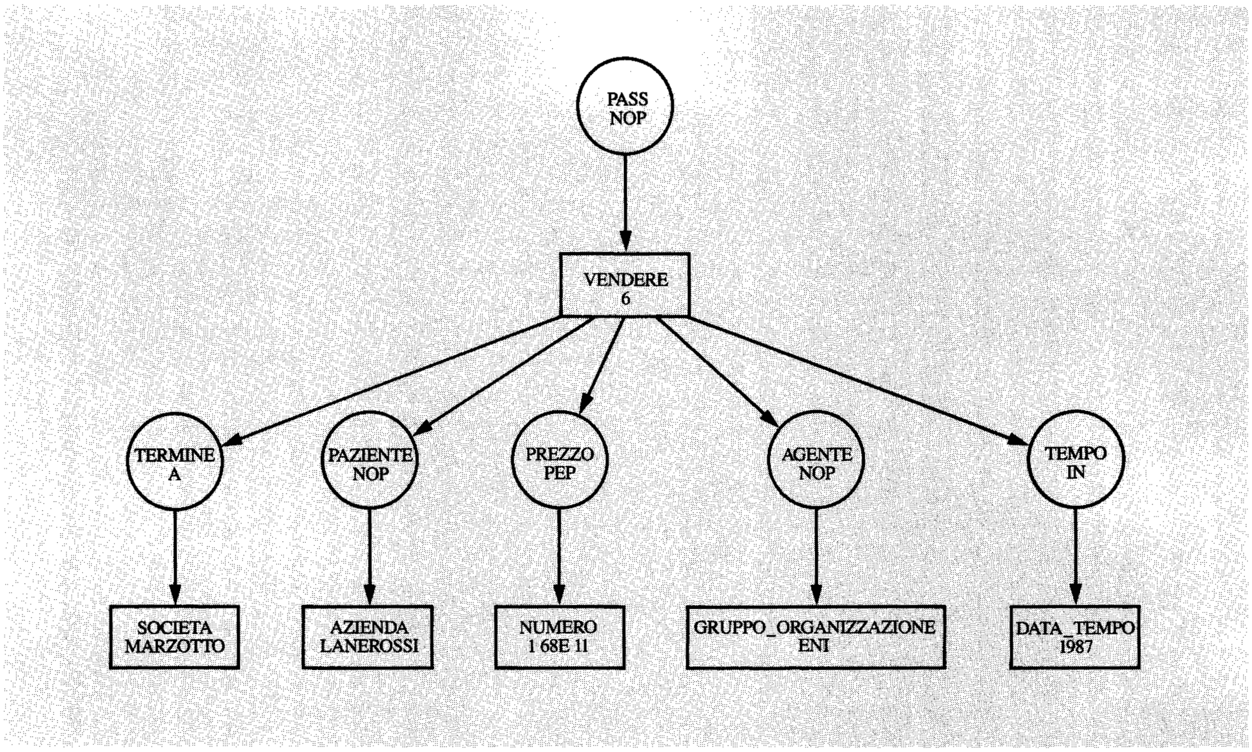

Figure 2

Conceptual graph for the sentence "Nel 1987 il gruppo Eni ha venduto alla Marzotto la Lanerossi per 168 miliardi." [*In 1987 the Eni group sold Lanerossi to Marzotto for 168 thousand million lire.*]

system is unable to answer questions such as *Who did buy Lanerossi? For how many years has Marzotto owned Lanerossi? Who was the previous owner of Lanerossi?* From this point of view we can say that the original System N was able to represent only the explicit information in a text, and to reach this goal it used a surface semantics represented by a semantic graph. The new work reported in this paper extends that information by incorporating more background knowledge, representing the deeper semantics.

Events and rules

The inference process is based on a set of semantic rules which are triggered by certain elements in the semantic representation of a sentence. We began with the assumption that the information that may be deduced from a text is often strictly connected to the type of event specified in the same text. By *event* we mean any verb which constitutes an expression "making sense," together with the concepts to which it refers (in "meet to discuss," *to discuss* is not an event).

We then defined two sets of rules. The first set contains rules whose application depends only on the kind of event in the analyzed text; the second, instead, contains rules which make inferences only from the conceptual relations appearing in the graphs, independently of the kind of event. The input for each rule is a conceptual graph containing just one event, and the task of the rule is to record the inferred information in another conceptual graph.

We have restricted the number of input events to avoid redundancies. In fact, if a rule had as input a graph containing more than one event, with only one being meaningful for the deduction, we would find in the final graph not only the inferred information, but also other information concerning the events rejected by the rule. Because this information is already recorded in the first graph, we would have a redundancy of information.

For example, in the sentence

Il presidente approva il contratto e nomina segretari Carlo Rossi.
[*The president approves the contract and appoints Carlo Rossi secretary.*]

there are two events described by the verbs *approvare* [*to approve*] and *nominare* [*to appoint*]. The semantic rule concerning *to appoint* is able to infer that if a person has been appointed secretary, *that person is a secretary.* This kind of deduction can seem trivial, but it allows users to obtain an answer to the questions *Who is a secretary? Who is Carlo Rossi?*

The inferred information is recorded in the form of a conceptual graph. It would be useless for that graph to contain *The president approves the contract*, because that information is already recorded in the first graph, obtained without any inference.

The semantic rules are not tied (correlated) to specific words describing an event, such as the verb *to appoint*; as shown in the following section, they depend more generally on classes of events.

Inference depending on the event

There are four classes of events that can trigger the deduction of implicit information. The first three classes belong to the conceptual hierarchy of system N:

- *Atti commerciali* [*commercial transactions*]: e.g., *vendere, acquistare* [*to sell, to buy*].
- *Atti elettivi* [*elective acts*]: e.g., *nominare, eleggere* [*to appoint, to elect*]. (Examples of these two classes were shown in the previous section.)
- *Atti decisionali* [*decision-making acts*]: e.g., *decidere, deliberare* [*to decide, to deliberate*].

For example, if we read that *Giovanni decide di andare al cinema* [John decides to go the cinema], we can infer that *Giovanni andrà al cinema* [John will go to the cinema].

The fourth class, *accomplishments*, represents a further subdivision in the hierarchy: It contains conclusive events (i.e., actions brought to a conclusion). In the hierarchy, this class is not only superior to the previous three classes, but also to the following classes:

- *Atti di communicazione* [*communication acts*]: e.g., *affermare, richiedere* [*to declare, to ask for*].
- *Atti terminativi* [*ending acts*]: e.g., *concludere, finire* [*to conclude, to end*].
- *Atti direttivi* [*management acts*]: e.g., *amministrare, presiedere* [*to manage, to chair*].
- *Atti distributivi* [*distributive acts*]: e.g., *conglobare, unire* [*to merge, to join*].
- *Atti effettivi* [*effective acts*]: e.g., *coprire, realizzare* [*to cover, to realize*].
- *Atti induttivi* [*inductive acts*]: e.g., *coinvolgere, soddisfare* [*to involve, to satisfy*].
- *Atti partecipativi participative acts*]: e.g., *partecipare, riunire* [*to participate, to get together, to meet*].
- *Atti modificatori di possesso* [*possession-modifying acts*]: e.g., *date, perdere* [*to give, to lose*].

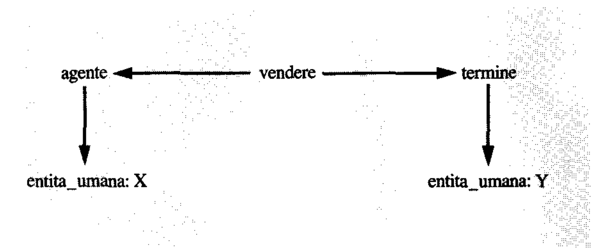

Figure 3

Generalized conceptual graph for *entità umana X vende ad entità umana Y* [*human entity X sells to human entity Y*].

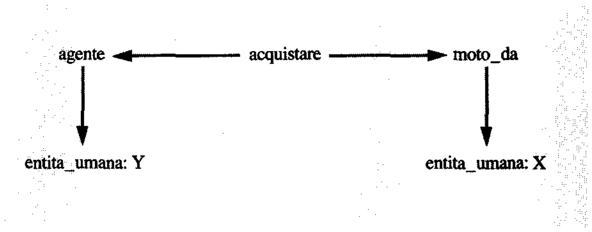

Figure 4

Generalized conceptual graph for *entità umana Y acquista da entità umana X* [*human entity Y buys from human entity X*].

- *Atti modificatori di caratteristica* [*characteristic-changing acts*]: e.g., *degradare, rovinare* [*to degrade, to ruin*].

In this paper, we consider only the rules concerning *commercial transactions* and *accomplishments.*

- *Commercial transactions*

Two kinds of rules are triggered by these events. The purpose of the first is to reverse, if possible, the action (e.g., *sell* to *buy* and vice versa). The second one works on the structures implying possession. The first rule can be described as follows:

entità umana X vende ad entità umana Y ↔ *entità umana Y acquista da entità umana X*
[*human entity X sells to human entity Y* ↔ *human entity Y buys from human entity X*]

where *human entity* indicates a class of the conceptual hierarchy and the double arrow means that the rule can be read from left to right or vice versa. We suppose that the input event is *to sell.* The rule is applied if the structure in **Figure 3** is recognized. If so, the graph, **Figure 4**, is

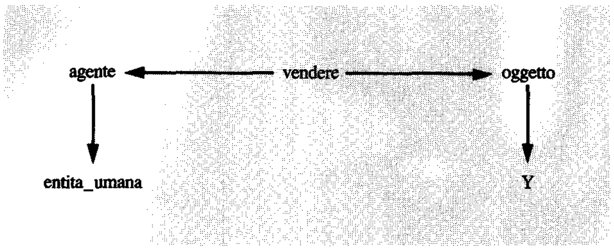

Figure 5

Generalized conceptual graph for *entità umana vende l'oggetto Y* [*human entity sells object Y*].

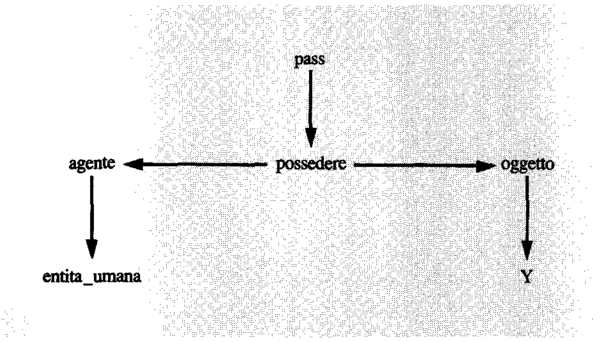

Figure 6

Generalized conceptual graph for *entità umana possedeva l'oggetto Y* [*human entity owned the object Y*].

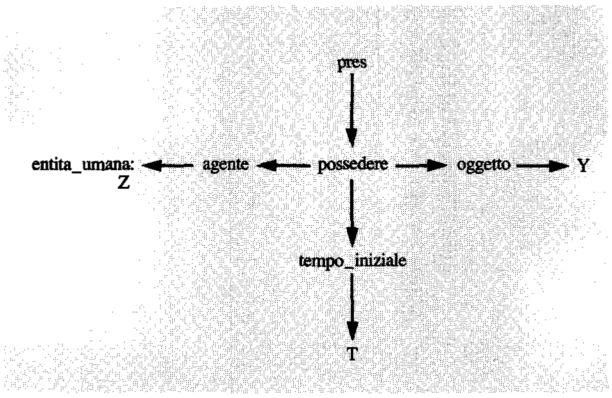

Figure 7

Generalized conceptual graph for *entità umana Z possiede l'oggetto Y dal tempo T* [*human entity Z has owned object Y since time T*].

produced, in which the event and a conceptual relation are changed. The same operations can be repeated, in reverse order, if the event in the graph represents a purchase.

This type of event also implies the concept of ownership:

entità umana vende l'oggetto Y
[human entity sells object Y]

This structure is expressed in the form of a generalized conceptual graph in **Figure 5**. If this hypothesis is verified, the graph is modified as shown in **Figure 6**.

Often in this kind of sentence, we find time complements conceptually related to the event. For example,

entità umana X vende ad entità umana Z l'oggetto Y al tempo T
[human entity X sells to human entity Z object Y at time T]

In this case the rule produces the inferred graph of **Figure 7** corresponding to the expression

entità umana Z possiede l'oggetto Y dal tempo T
[human entity Z has owned object Y since time T]

As a complete example, let us consider again the previously analyzed sentence

Nel 1987 il gruppo Eni ha venduto alla Marzotto la Lanerossi per 168 miliardi.
[In 1987 the Eni group sold Lanerossi to Marzotto for 168 thousand million lire.]

Both rules produce three inferred graphs, corresponding to the following expressions:

Nel 1987 la Marzotto ha acquistato la Lanerossi dal gruppo Eni per 168 miliardi.
[In 1987 Marzotto bought Lanerossi from the Eni group for 168 thousand million lire.]
Il gruppo Eni possedeva la Lanerossi.
[The Eni group owned Lanerossi.]
La Marzotto possiede la Lanerossi dal 1987.
[Marzotto has owned Lanerossi since 1987.]

- *Accomplishments*

As we have noted previously, this class describes conclusive events. The semantic rule for this category is more general than the one for transactions, because it includes a large number of events. The deduction is based on the assumption that if an accomplishment happens for a given purpose, then the purpose itself also happens (*to convene* **in order to vote** also means *to vote*, while **to change the regulations in order to vote** does not necessarily mean *to vote*, because *to change* is not an accomplishment under our definition).

In the case for which the accomplishment has no purpose, the rule verifies whether the "event" has an argument, whose purpose is represented by a concept, such that a verb (an action) is derivable from it. For example, *firmare il contratto di acquisto* [*to sign a purchase agreement*] can be referred to the verb *acquistare* [*to purchase*]. The concept *acquisto* (*the purchase*) in the clause represents the purpose of the agreement, which is the argument of *to sign*, but not the purpose of the event, as is shown in the conceptual graph of **Figure 8.**

In both cases considered by the rule, the purpose can be a noun (*purchase*) or a verb (*to vote*). In the first situation the system consults a table of semantic relationships between nouns and verbs to find a verb with the same meaning as the concept noun. When the verb has been found, it replaces the event, after the removal of the relationships of purpose and argument.

An example of the application of this rule is given by the sentence

La Ferruzzi firma il contratto di acquisto della Cifa.
[*Ferruzzi signs the purchase agreement for Cifa.*]

whose graph is shown in **Figure 9**. The application of the rule produces the graph in **Figure 10** after the replacement of the sequence *firmare il contratto di acquisto* [*to sign for the purchase agreement*] with the verb *acquistare* [*to purchase*], produced through semantic affinity with the noun *acquisto* [*purchase*].

The relevance of such a rule is to make explicit an event which, instead, is hidden in the context. This new event, belonging to one of the classes that trigger other inferences (e.g., *commercial transactions*), recalls other deduction rules.

In fact, the inferred graph of the previous example is filtered by the rule that takes commercial transactions into account. This rule produces another inferred graph:

La Ferruzzi possiede la Cifa.
[*Ferruzzi owns Cifa.*]

Inference independent of the event

The second set of rules does not depend on the kind of event contained in a sentence. These rules modify certain conceptual relations in the conceptual graph produced by the surface semantics of system N, and try to deduce other possible graphs. Let us consider the sentence

L'amministratore delegato approva con gli azionisti il bilancio.
[*The managing director approves the balance with the shareholders.*]

To the question *Chi approva il bilancio?* [*Who approves the balance?*], System N answers that the managing director approves. Instead, the more correct answer is that both the managing director and the shareholders approve.

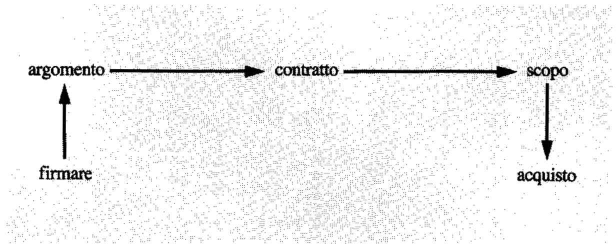

Figure 8

Conceptual graph for *firmare il contratto di acquisto* [*to sign a purchase agreement*].

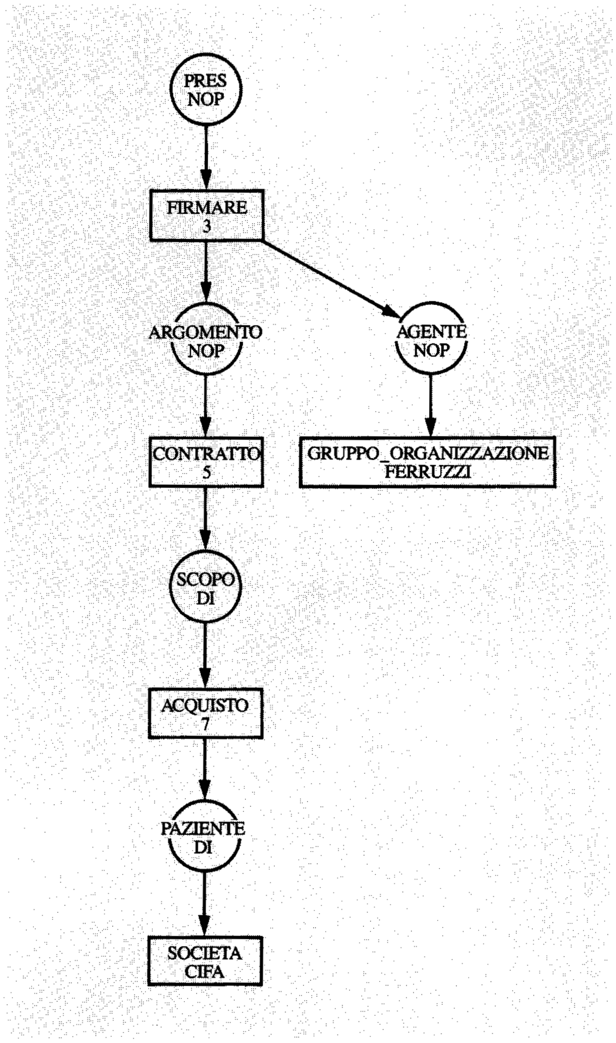

Figure 9

Conceptual graph for the sentence *La Ferruzzi firma il contratto di acquisto della Cifa* [*Ferruzzi signs the purchase agreement for Cifa*].

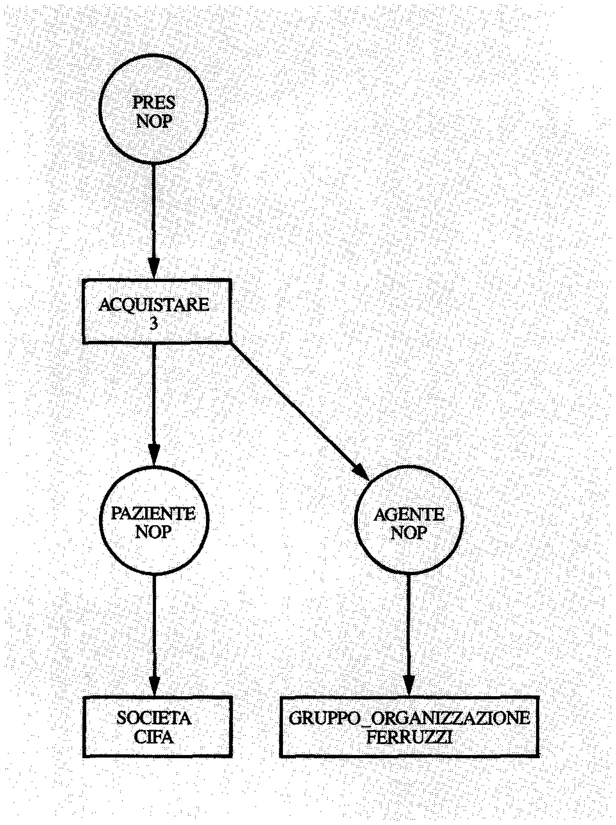

Figure 10

Conceptual graph for the sentence *La Ferruzzi acquista la Cifa* [*Ferruzzi acquires Cifa*].

System N gives the first answer because it looks only for the agents of the action, and *shareholders* is not an agent. Such a problem can easily be solved by creating a new graph in which the concepts *managing director* and *shareholders* exchange their roles. The rule is triggered by the existence of the relationship of "company" in the conceptual graph. This relationship is considered as an "agent" relationship of the same event, and the other semantic relationships of the original graph are true also for the deduced graph.

The inference process

The inference process begins during the phase of text understanding, and works through the following phases:

- System N analyzes a sentence and produces a surface conceptual graph G.
- G is analyzed to determine whether its event belongs to the previously considered hierarchical classes. In such a case all the conceptual relations linking the event to the nearby words are identified. If more events appear, the graph is divided into as many parts as there are events. This is realized through a recursive process that builds, for each event a, a graph G_a containing not only this kind of triplet:

  [*a*] ← (*conceptual_relation*) ← [*concept_in*]

  and this other:

  [*a*] → (*conceptual_relation*) → [*concept_out*],

  but also this:

  [*c_in*] → (*conceptual_relation*) → [*c_out*],

  which verifies one, and only one, of the following properties:
  1. *c_in* is the concept *going out* from one of the conceptual relations of G_a, and the graph [*c_in*] → (*conceptual_relation*) → [*c_out*] is contained in G but not in G_a.
  2. *c_out* is the concept *going in* to one of the conceptual relations of G_a, and the graph [*c_in*] → (*conceptual_relation*) → [*c_out*] is contained in G but not in G_a.
- Any single graph G_a produced in the previous phase is examined by the semantic rules, after a timely analysis of the event a. If the event is not significant in the determination of information, then the only activated rules are those that do not consider the kind of the current action. In this way a set $IG_a = \{G_a^1, \cdots, G_a^n\}$ is produced, where G_a^i represents the ith graph inferred by G_a. This set could be empty, and this happens when none of the rules succeeds.
- If IG_a is a nonempty set, the inference process examines, for each element G_a^i of IG_a, the event a_i. When $a_i \neq a$, the process begins again in order to find out further information. This generally happens when the event a is a commercial, decision-making act or an accomplishment; these are, in fact, the rules that modify the action. For each graph G_a^i a new set $IG_a^i = \{G_a^{i_1}, \cdots, G_a^{i_m}\}$ is produced, where $G_a^{i_k}$ represents the kth graph inferred by G_a^i. This recursive process ends when the application of the rules to the graphs belonging to the previous level fails. The inference process may end at the jth level for either of two different reasons:
  1. The action of the graphs at level j is equal to the action of the graphs at level $j - 1$.
  2. The action of the graphs at level j is different from the action of the graphs at level $j - 1$, but the semantic hypothesis of the inference rules fails.
- All the inferred graphs are recorded in the knowledge base.

This process increases the capability of System N to answer questions, in Italian, concerning the deep meaning of the analyzed texts.

Uncertainty in the text-understanding system

The general problem of the deduction of implicit information can be partially solved by means of semantic rules, but one of the specific problems that is difficult to solve is the processing of imprecise questions.

Let us suppose that the knowledge base contains the semantic representation for the sentence

La Fiat Auto ha raggiunto nel 1987 una quota del mercato automobilistico del 14,3%.
[*In 1987 Fiat Auto attained a 14.3% share of the car market.*]

Among the possible questions allowed by System N, it would be interesting and useful to ask

La Fiat Auto ha raggiunto una grossa quota del mercato automobilistico?
[*Has Fiat Auto attained a large share of the car market?*]

At present, it is not possible for System N to answer this question. The quantity relation in the conceptual graphs, whether of the declarative sentence or of the questions, links different concepts in the two representations, as shown in the corresponding subgraphs for the sentence and for the question, respectively:

[*number_14.3*] ← (*quantity*) ← [*car_market*]
[*large*] ← (*quantity*) ← [*car_market*]

Thus, the match between the semantic representations of the sentence and the question, to find a correct answer graph, fails.

Managing the problem of uncertainty is not trivial in a text-understanding system in which the knowledge representation is obtained by means of a precise formal structure such as conceptual graphs. The use of inference rules is not an efficient approach, because it is not possible to previously classify the kinds of imprecise questions that a user can formulate.

The main problem with such questions is that generally the match between them and the knowledge base fails, even when it is really possible to extract satisfiable answers. To get an answer to this problem, we are working in two directions: a correct interpretation of the concepts obscured by vagueness, and an alternative way to perform the matching between the questions and the sentences of the knowledge base. At the moment we are dealing with the problem but using the formalism of fuzzy logic [15], both in the data representation and in the inference rules [4].

According to the contents of the knowledge base, we have considered questions containing

- Quantity imprecision (i.e., *a lot, a few*)

Il bilancio della Rambaudi é stato di molti miliardi?
[*Was Rambaudi's budget many thousand millions of lire?*]

- Time imprecision (i.e., *at the beginning of the year, in many years*)

Che cosa ha acquistato la Montedison recentemente?
[*What did Montedison recently buy?*]

The set of questions containing quantity or time imprecisions has been subdivided into two subsets, Boolean and non-Boolean. By Boolean questions, we mean interrogative utterances without an interrogative pronoun; that is, questions that require an affirmative or negative answer:

La Ferruzzi ha speso molto per l'acquisto della Cifa?
[*Did Ferruzzi spend a lot on the Cifa purchase?*].

Non-Boolean questions represent utterances that contain an interrogative pronoun:

Cosa ha acquistato l'Ansaldo recentemente?
[*What did Ansaldo recently buy?*]

Boolean questions

A specific example will be used to show the processing of Boolean questions, through the formalism of fuzzy logic. Let us consider the sentence

La Ferruzzi ha speso 39 miliardi per l'acquisto della Cifa.
[*Ferruzzi spent 39 thousand million lire on the Cifa purchase.*]

A possible question could be

La Ferruzzi ha speso molto per l'acquisto della Cifa?
[*Did Ferruzzi spend a lot on the Cifa purchase?*]

The conceptual graph for the declarative sentence is the following:

[*COMPANY:FERRUZZI*] ← (*AGENT*) ← [*TO SPEND*]

[*PURCHASE*] ← (*GOAL*) ← [*TO SPEND*]

[*COMPANY:CIFA*] ← (*PATIENT*) ← [*PURCHASE*]

[NUMBER_14.3%] ← (*QUANTITY*) ← [*TO SPEND*]

while the conceptual graph for the question is the following:

[*COMPANY:FERRUZZI*] ← (*AGENT*) ← [*TO SPEND*]

[*PURCHASE*] ← (*GOAL*) ← [*TO SPEND*]

[*COMPANY:CIFA*] ← (*PATIENT*) ← [*PURCHASE*]

[*A_LOT*] ← (*QUANTITY*) ← [*TO SPEND*]

The representations fail to match, because the question graph is not completely contained in the sentence graph.

In this case we eliminate temporarily the time and quantity relations from the question graph and then try to match the representations; i.e., we perform a "partial matching." The result is that now the system is able to select, in the knowledge base, the graph from which the answer can be extracted. When the quantity relation is ignored, the question graph becomes

$[COMPANY{:}FERRUZZI] \leftarrow (AGENT) \leftarrow [TO\ SPEND]$

$[PURCHASE] \leftarrow (GOAL) \leftarrow [TO\ SPEND]$

$[COMPANY{:}CIFA] \leftarrow (PATIENT) \leftarrow [PURCHASE]$

The graph matching is now successful. This process does not yet give the right answer, however, so the next step is to define, according to the fuzzy logic formalism [16], a universe of discourse U, which is constituted by integer numbers in the set [0, 100] [the numbers represent the thousands of millions (billions) spent by a company]. We also consider a linguistic variable *cost* and its term set, for example

$$T(cost) = A\_LOT + VERY\_LOT + A\_FEW + NOT\_A\_FEW + NOT\_A\_LOT + \cdots$$

We can assume A_LOT as the primary term and represent it with the fuzzy set

$$A\_LOT = 1/(80 - 100) + 0.8/(60 - 79) + 0.4/(40 - 59),$$

in which $\mu/(\alpha - \beta)$ means that all the integers between α and β have a value μ of membership to the set. Starting from the primary term A_LOT, with μ_{A_LOT} as membership function, we can build the fuzzy set A_FEW in the same way that *false* is derived from *true*,

$$\mu_{A\_FEW}(u) = \mu_{A\_LOT}(100 - u), \qquad u \in U.$$

From this follows

$$A\_FEW = 0.4/(41 - 60) + 0.8/(21 - 40) + 1/(0 - 20).$$

Given a fuzzy set A and its membership function μ_A, the operators *very, more or less, not, and, or* are defined in the following way, where u represents an element of a universe of discourse U:

$$\mu_{veryA}(u) = \mu_A(u)^2,$$

$$\mu_{more \vee less A}(u) = \mu_A(u)^{1/2},$$

$$\mu_{notA} = 1 - \mu_A(u),$$

$$\mu_{A \wedge B} = \mu_A(u) \wedge \mu_B(u)$$

($\wedge$ represents the minimum between two elements), and

$$\mu_{A \vee B} = \mu_A(u) \vee \mu_B(u)$$

($\vee$ represents the maximum between two elements).

By the application of these operators, we can build the other terms of $T(cost)$ by means of the assumption

$$VERY\_LOT = 1/(80 - 100) + 0.64/(60 - 79) + 0.16/(40 - 59);$$

$$NOT\_A\_LOT = not\ A\_LOT = 0.2/(60 - 79) + 0.6/(40 - 59) + 1/(0 - 39);$$

$$NOT\_A\_FEW = not\ A\_FEW + 0.6/(41 - 60) + 0.2/(21 - 40);$$

$$VERY\_FEW = very\ A\_FEW = 0.16/(41 - 60) + 0.64/(21 - 40) + 1/(0 - 20);$$

etc.

By considering the previous question

La Ferruzzi ha speso molto per l'acquisto della Cifa?
[*Did Ferruzzi spend a lot on the Cifa purchase?*]

and by using this formalism, we can translate it into the proposition

$$cost\ is\ A\_LOT. \tag{1}$$

We assume that the sentence

La Ferruzzi ha speso 39 miliardi per l'acquisto della Cifa.
[*Ferruzzi spent 39 thousand million lire on the Cifa purchase.*]

has been selected by the system through the partial match. We can find, among the fuzzy sets belonging to $T(cost)$, the fuzzy set in which 39 has the maximum membership value. If this happens in more than one set, we can select one of them by means of the following hierarchy:

1. Primary terms.
2. Terms built from others through the operators *very, more or less, not.*
3. Terms built from others through the operator *and.*
4. Terms built from others through the operator *or.*

Referring to our example, we find the fuzzy set NOT_A_LOT, so we can translate the sentence into the proposition

$$cost\ is\ NOT\_A\_LOT. \tag{2}$$

On propositions (1) and (2) we can apply the rules of *implicit conjunction* and *maximal restriction* [17]: If P and Q are fuzzy sets of universes U and V respectively, the rule of implicit conjunction asserts that from the propositions x *is* P and y *is* Q follows

$$x\ is\ P \wedge y\ is\ Q. \tag{I}$$

In the same hypothesis, from the propositions x *is* P and x *is* Q for the rule of maximal restriction follows

$$x\ is\ P \cap Q, \tag{II}$$

where $P \cap Q$ has the following membership function:

$$file\ \mu_{P \cap Q}(u) = \mu_P(u) \wedge \mu_Q(u), \qquad u \in U.$$

The application of (I) and (II) to propositions (1) and (2) allows the following propositions to be generated:

cost is S, (3)

where $S = A\_LOT \cap NOT\_A\_LOT$.

If S belongs to the term set of the variable *cost*, we have found an answer. Instead, if S does not belong to $T(cost)$, it can be linguistically approximated by means of an element of $T(cost)$ [18]. This process can be summarized as follows:

1. A question Q is given in input. The system executes the partial matching between its interpretations and the knowledge base and selects a sentence S from which the answer can be derived.
2. The system builds a universe of discourse U, a linguistic variable u referring to the imprecise concept c contained in Q, and its term set $T(u)$.
3. From Q and S the system builds the propositions *u is P* and *u is R*, where P is the fuzzy set of U having as its label the imprecise concept c of Q, and R is the fuzzy set of U in which the term of the sentence referring to c has the maximum membership value.
4. By means of the rules of implicit conjunction and maximal restriction, the system infers the proposition *u is* $P \cap R$. If $P \cap R$ or its linguistic approximation is equal to P, the answer is affirmative; else the answer is negative.

This simple reasoning can easily be extended to questions containing more than one imprecise concept.

- *Non-Boolean questions*

This class contains questions in which there is either an imprecise concept or a question pronoun. For example,

Cosa ha acquistato la Fiat nella prima metà dell'anno?
[*What did Fiat buy in the first part of the year?*]

It is impossible to handle this kind of question by means of the process for Boolean questions. In fact, from this question we cannot deduce the propositions of fuzzy logic. We suppose that the knowledge base contains the semantic graphs for the following sentences:

- *Nel 1980 la Montedison ha firmato un accordo con la Rodriguez.*
- [*In 1980 Montedison signed an agreement with Rodriguez.*]
- *Nel giugno 1987 la Montedison ha acquistato il 15% della Rio.*
- [*In June 1987 Montedison bought 15% of Rio.*]
- *Nel 1987 la Montedison ha proposto la creazione di un polo chimico nazionale.*
- [*In 1987 Montedison suggested the foundation of a national chemical pool.*]

Table 1 Membership values for elements of the fuzzy set *RECENT* and its derivative sets.

	REC	*VREC*	*NREC*	*NVREC*
89	1	1	0	0
88	1	1	0	0
87	0.9	0.81	0.1	0.19
86	0.7	0.49	0.3	0.51
85	0.5	0.25	0.5	0.75
84	0.4	0.16	0.6	0.84
83	0.3	0.09	0.7	0.91
82	0.2	0.04	0.8	0.96
81	0	0	1	1
80	0	0	1	1
79	0	0	1	1
78	0	0	1	1
77	0	0	1	1
76	0	0	1	1
75	0	0	1	1

We now ask the question

Cosa ha fatto la Montedison recentemente?
[*What about Montedison recently?*]

To obtain an answer, we can build a universe of discourse $U = \{1975, 1976, \cdots, 1989\}$ and a linguistic variable *time* on U. The term set of *time* can be, for example,

$$T\,(\text{time}) = RECENT + VERY\_RECENT + NOT\_RECENT + \cdots.$$

We choose *RECENT* as the primary term. The other terms can be derived from it through the operators *very, more or less, and, or, not.* If μ_{REC} is the membership function of *RECENT*, the membership functions of the other terms will be the following:

$$\mu_{VERY\_RECENT}(u) = \mu_{very\ RECENT}(u) = \mu^2_{REC}(u), \qquad u \in U;$$

$$\mu_{NOT\_RECENT}(u) = \mu_{not\ RECENT}(u) = \mu^{1/2}_{REC}(u), \qquad u \in U; \text{ etc.}$$

Now we suppose that the fuzzy set *RECENT* is represented thus:

$$RECENT = 1/89 + 1/88 + 0.9/87 + 0.7/86 + 0.5/85 + 0.4/84 + 0.3/83 + 0.2/82.$$

The situation can be summarized by **Table 1**, in which *REC, VREC, NREC, NVREC* represent respectively the fuzzy sets *RECENT, VERY_RECENT, NOT_RECENT, NOT_VERY_RECENT*, and the numbers represent the membership value of the elements to the described fuzzy sets.

As we can see from the table, the years with the maximum value of membership to the fuzzy set *RECENT* are 1988 and 1989. We can then obtain from the knowledge base, by means of standard matching, an answer for the questions

- *Cosa ha fatto la Montedison nel 1989?*
- *[What about Montedison in 1989?]*
- *Cosa ha fatto la Montedison nel 1988?*
- *[What about Montedison in 1988?]*

In the same way, if the question is

Cosa ha fatto la Montedison non molto recentemente?
[What about Montedison not very recently?]

we can find out from the knowledge base something about Montedison from 1975 to 1981.

Conclusions and future developments

Processes which seem to require only a minimum of "intelligence" almost always represent complex problems. What we really want is for the computer to understand what we say. To approach this goal, a computer must read sentences and make inferences about likely circumstances, presuppositions, and conclusions. The work is tedious and difficult. We have begun to extend a superficial text-understanding system, providing it with some inferential capability. At the moment the semantic inference rules we have developed cover a set of typical events in the context of a domain based on economics and finance. However, the rules are also valid in different contexts, and the generality of the mechanism based on the proposed approach permits the development of new inference rules about events.

Users often formulate imprecise questions requiring precise answers. To solve such problems, we have designed an algorithm based on fuzzy logic which will tolerate questions containing imprecise terms. The algorithm is currently being implemented and tested. System N and the inference process have been implemented on IBM System/370™ architecture, under VM/CMS, using VM/Prolog as the host language.

System/370 is a trademark of International Business Machines Corporation.

References

1. M. Minsky, "A Framework for Representing Knowledge," *The Psychology of Computer Vision*, P. H. Winston, Ed., McGraw-Hill Book Co., Inc., New York, 1979.
2. R. C. Schank and R. P. Abelson, *Script, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structure*, Lea Lawrence Erlbaum Association, Hillsdale, NJ, 1977.
3. K. Dalhgreen and J. McDowell, "Commonsense Reasoning with Verbs," *Proceedings of the International Joint Conference on Artificial Intelligence*, Milan, 1987, pp. 446–448.
4. C. M. Calamani, "Deduzione di Informazione Implicita nell'Analisi del Linguaggio Naturale," graduation thesis in mathematics, University of Rome "La Sapienza," Rome, Italy, 1988/89.
5. F. Antonacci, M. T. Pazienza, M. Russo, and P. Velardi, "Representation and Control Strategies for Large Knowledge Domains: An Application to Natural Language Processing," *J. Appl. Artif. Intell.* **2**, 213–249 (1988).
6. F. Antonacci, M. T. Pazienza, M. Russo, and P. Velardi, "A System for Text Analysis and Lexical Knowledge Acquisition," *Data & Knowledge Eng.* **4**, 123–127 (1989).
7. F. Antonacci, S. Magrini, and C. Ronchini, "An Approach to Intelligent Information Retrieval Bridging Text Search and Text Understanding," *Proceedings of the Second I.T.L. Conference on Natural Language Processing*, Paris, March 13–15, 1990, pp. 197–218.
8. M. Russo, "A Generative Grammar Approach for the Morphologic and Morphosyntactic Analysis of Italian," *Proceedings of the Third Conference of the European Chapter of the ACL*, Copenhagen, April 1–3, 1987, pp. 32–37.
9. M. Russo, "A Rule Based System for the Morphologic and Morphosyntactic Analysis of the Italian Language," *Natural Language Understanding and Logic Programming II*, V. Dahl and P. Saint-Dizier, Eds., North-Holland Publishing Co., Amsterdam, 1988, pp. 225–237.
10. F. Antonacci, "Using Semantic Hints in a Syntactic Parser," *Technical Report SCR-0004*, IBM Scientific Center, Rome, Italy, May 1987.
11. F. Antonacci, P. Velardi, and M. T. Pazienza, "A High Coverage Grammar for the Italian Language," *J. Assoc. Literacy & Linguist. Computing* **4**, 210–214 (1989).
12. M. T. Pazienza and P. Velardi, "A Structured Representation of Word Senses for Semantic Analysis," *Proceedings of the Third Conference of the European Chapter of the ACL*, Copenhagen, April 1–3, 1987, pp. 249–257.
13. J. F. Sowa, *Conceptual Structures: Information Processing in Mind and Machine*, Addison-Wesley Publishing Co., Reading, MA, 1984.
14. P. Velardi, M. T. Pazienza, and S. Magrini, "Extensive Acquisition of a Semantic Lexicon," *Proceedings of the Workshop on Conceptual Graphs*, American Association for Artificial Intelligence, St. Paul, MN, August 27, 1988, pp. 278–285.
15. H. J. Zimmermann, *Fuzzy Set Theory and Its Application*, Kluwer–Nijoff Publishers, London, 1985.
16. L. A. Zadeh, "Fuzzy Logic and Its Application to Approximate Reasoning," *Info. Proc.* **10**, 34–39 (1974).
17. *Fuzzy Sets and Their Application to Cognitive and Decision Process*, L. A. Zadeh, K. Fu, S. Tanaka, and K. Shimura, Eds., Academic Press, Inc., New York, 1975.
18. E. H. Mamdani and B. R. Gaines, *Fuzzy Reasoning and Its Application*, Academic Press, Inc., New York, 1981.

Received November 18, 1990; accepted for publication January 18, 1992

Francesco Antonacci *IBM SEMEA Scientific and Technical Solution Centre, viale O. Pacifico 171/173, 00144 Rome, Italy (ANTONACC at ROMESC, antonacc@romesc.vnet.ibm.com).* Mr. Antonacci is a Research Staff Member in the Italian Language Technology group. After receiving the B.S. degree in chemistry in 1972 from the University of Bari, Italy, he worked as a research chemist in the Agricultural Laboratory of the National Council of Research. In 1974 he joined the Scientific Centers of IBM, working on the design and development of a relational database and a formal query language to access it. Subsequently, he contributed to the development of a geographic database management system and its applications. From 1985 until 1989 he worked in the "Knowledge Representation" group on the development of an Italian text understanding system and techniques of intelligent information retrieval. Beginning in 1990, in the Italian Language Technology group, Mr. Antonacci contributed to the development of a product, RI/VM, used for intelligent information retrieval, and of applications in education to teach the Italian language.

Cecilia M. Calamani *TECSIEL S.P.A., v. del Ponte di Piscina Cupa 43, 00129 Rome, Italy.* Dr. Calamani received the B.S. degree in mathematics from the University of Rome in 1989; in conjunction with her academic work, she developed a thesis on inference and deduction of natural language at the IBM Scientific Center in Rome. From 1989 until 1991 she worked at the Computer Center of Rome University "La Sapienza" on natural-language understanding. In 1991 Dr. Calamani joined TECSIEL S.P.A.