

INTERPRETATION OF NOUN PHRASES IN INTENSIONAL CONTEXTS

Leonardo LESMO and Paolo TEREZIANI

Dipartimento di Informatica
Universita' di Torino
C.so Svizzera 185 - 10149 TORINO - ITALY

ABSTRACT

This paper presents a network formalism for representing the meaning of noun phrases occurring in the context of intensional verbs such as *seek* and *want*. The basic assumption is that an intermediate representation is useful to carry out the interpretation process perspicuously. The proposed representation keeps apart *de-re* (transparent) and *de-dicto* readings, stating, by means of suitable arcs and nodes, that in the first case there is a real *counterpart* of the described entity, whereas in the second case no such counterpart exists. Moreover, the concept of *relevance* of the description is emphasized and the relationships between intensional contexts and the dicotomy existing between *value-free* and *value-loaded* readings of definite descriptions is discussed. Finally, the impact of the intermediate representation on the discourse history is considered, in order to explain how such a representation accounts for the contextual data and updates them according to the informational contents of the sentence being analyzed.

1. INTRODUCTION

Formally, an intension is a mapping that associates with every possible world (actually, with indices that identify possible worlds) a denotation, where a denotation is a set of entities (individuals, relations, functions). If we take a term, say "animal", it *denotes*, in the present world, a set of individuals; but the meaning of the term cannot be characterized by that set. For instance, the terms "human" and the term "featherless biped" *denote* the same set, but their *meaning* is different, since the fact that there are no featherless bipeds except human beings, in this world, can be regarded as a mere accident. It is not difficult to conceive of a different world where another species shares the same property with humans.

The characterization of possible worlds has been subject to deep philosophical debates, which were mainly concerned with the relationships existing between "conceivable worlds" and "possible worlds", as well as with the connections existing between alternative worlds, but, for the present purpose, we only need to note that:

- The concept of "possible world" is rather easy to render formally, expanding the usual concept of *interpretation* in logics by introducing "indices" on worlds. Thus, for instance, the mapping which assigns to each predicate symbol a relation (of the same arity) on the universe, must be viewed as a set of mappings (one for each world), since the *denotation* of the predicate can vary from world to world. Equivalently, the *intension* of a n-ary predicate can be defined as a function from world indices to n-ary relations.

- The concept of "possible world" is rather common in the everyday practice. For instance, in saying

(1) John seeks a unicorn

no relation between *John* and a real world entity referred to by "a unicorn" can be assumed, since no such entity exists. Nonetheless

(2) John seeks a centaur

is different from (1). Intuitively, the expression "centaur" has a meaning different from "unicorn", because in some (fictitious)

world they are used to refer to different entities. The same cannot be said for expressions such as "centaur" and "a being half human and half horse", since the second can be used to define the first¹. On the other hand, the fact that possible worlds are common to logics and intuition does not mean that there is an exact match in the way the concept is used.

Let's consider, for instance, the best known treatment of natural language in logics, i.e. the work of Montague's. In order to maintain the strict compositionality of his approach, Montague had to make the interpretation of linguistic entities considerably complex². For instance, the interpretation of a NP is given as a set of properties, where a *property* is a function from indices to set of individuals. For instance the *property* "redness" is that function that, for each index, says what is red at that index. Of course, this is intuitively acceptable, but why should "John" be interpreted as a set of properties? It is true that with any person we know we associate a bundle of properties that that person has (i.e. *properties* such that the individual denoted by "John" happens to be in their denotation evaluated at this index), but it seems also true that when I say "John" I refer to an individual and not to his properties³. Even more complex is the treatment of NP's in object position; for instance the phrase "a unicorn" in (1) is interpreted as a *property of properties*, in order to account for the fact that, loosely speaking, sets have to be evaluated "here and now" and so the *set* of properties that a unicorn has is the same as the *set* of properties that a centaur has.

We cannot go in further details here (more detailed information can be found in [Dowty et al. 81]), but a couple of questions naturally arises: Is it possible to maintain the basic principles of a possible worlds analysis without introducing the complexities of Montague's intensional logic? Is it possible to introduce an *intermediate level* in the semantic interpretation process, which preserves the idea of compositionality?

Recently, two interesting proposals concerning the semantics of NL have been put forward. It seems that they can give some suggestions about how to answer the questions above. We will briefly outline them here, but we must first notice that we do not aim at presenting the philosophical position of those authors, since the discussion that follows is strongly biased to the views we hold about the

¹ This can simply be due to my scarce knowledge about mythology. If there is a more precise characterization of centaurs, such that it is possible to conceive of a world where a being half human and half horse is not in the denotation of "centaur", then the same situation as "featherless biped" arises.

² We cannot discuss here our disagreement with a conception of syntax that assigns to a sentence as (1) two different *syntactic* structures in order to cope with the so-called *specific* and *non-specific* readings (the one where a particular unicorn is wanted vs. the one where any unicorn could serve John's purpose). We can only note that, in our view, syntax is mainly concerned with surface phenomena. Note also that in [Cooper 83] a Montague-based proposal is presented, that overcomes this difficulty.

³ Formally, the two alternatives are equivalent if we assume that no two individuals share the same "total" set of properties. For this reason, the solution proposed by Montague is perfect, if we assume, as he did, that no *intermediate representation* of the meaning of linguistic entities is ever built. In fact, the formulae of intensional logics are only tools to represent *constraints on possible worlds*, so that their complexities was not a matter at hand. As we will see, our position is different.

topic at hand.

The first proposal is "Situation semantics" by Barwise and Perry [83]. They view the meaning of a sentence as a relation between the *discourse situation* (who are the speaker, the addressee, etc.), the *connections of the speaker* (i.e. its relationships with the environment), and a *described situation* (the event, or "type of event", which is being described). What is fundamental, in our view, is the rejection of truth values as "senses" of sentences. This move accounts for the non-equivalence of attitude reports involving logically equivalent embedded sentences. For instance,

(3) John knows that two plus two is four
is not the same as

(4) John knows that the square of 37 is 1369

although both facts are necessary truths (with respects to the difference between *meaning* and *intension*, see also [Lewis 72]).

Also important is the exploitation of what is called *resource situation* in the interpretation of constituent of sentences. In the present concern, we can fix our attention on NP's. Let us consider, for example, the meaning of definite noun phrases:

"We can also think of d, c [the π]⁴ as giving a partial function from situations e to individuals a_σ :

d, c [the π] (e) is defined and = a_σ
iff

d, c [the π] a_σ, e " [Barwise & Perry 83, pp.149-150]

where the latter means that in the context where the expression is evaluated the "relational" meaning of π must single out exactly one individual. If we have available a particular resource situation, then we get an individual (a_σ), but if such a context is not available, then the interpretation of the NP is just the non-evaluated function from situations to individuals. In the first case, we get the *referential* (or *value-loaded*) use, since any predicate the sentence expresses is about the referred individual (a_σ); in the second case we get the *attributive* (or *value-free*) use; what is the predicate about in this case? It will be about the *eventual* (possible, future, fictional) filler of the description. The simplest way to obtain such a result is to incorporate the description of the individual into the interpretation of the sentence, so that, whenever the actual filler will be available, the predicate will be true of it. This approach does not seem to us in contrast with a possible worlds analysis: in each such a world there will be an individual denoted by the definite expression (possibly different individuals at different indices); thus, the *property* of satisfying the description will imply the *property* expressed by the predicate. In any case, the interpretation of sentences on the basis of some specific contextual situation suggests that the context is made explicit, in particular the *connections of the speaker*; of course, in an "understanding" system (be it a human or a computer) such connections are, at best, presumed, that is the hearer must "imagine" the connections on the basis of a "shared" environment or whatever she/he knows about the speaker. More important, in case the hearer is a machine, no real sharing of environment can be assumed, so that it must use an "internal" representation of the outer context; In the present work, this representation is available in terms of general world knowledge (represented in a KL-ONE like format), plus the history of the user-machine dialogue carried on so far.

Another piece of research that must be taken into account concerns *mental spaces* [Fauconnier 85]. This framework provides us with a tool for representing the relationships between the individuals mentioned in a sentence and their particular view of reality. We believe that this is one of the most common uses of the pre-theoretical notion of *possible world*. Of course, the knowledge, the beliefs, the opinions, the desires of different persons are not the same, and when we want to understand natural language we must try to model these differences. Although *mental spaces* are not claimed to be such a modelling technique⁵, nonetheless they seem useful in modelling the process that leads to the construction of the correct interpretation⁶.

⁴ "d" stands for the discourse situation, c for the "connections", "[the π]" for the meaning of 'the π ', "e" for the described situation.

This paper does not present an implementation of the *mental spaces* approach, but uses one of the basic ideas of that work: some linguistic contexts "set up" local spaces that constrain the meaning of the entities occurring in them. So, a mental space represents a local view of linguistic entities, that may be used to isolate the (referred) beliefs, desires, etc., of the mentioned individuals.

It should be clear from the above discussion that we are not adopting in full any of the positions of the reviewed works. It seems, however, that each of them can have a strong impact on the overall framework. Thus, an adaptation of mental spaces can help in modelling the process of building the meaning of sentences; a "situational" analysis can help in understanding how such a meaning must merge with the environment; a possible world approach can provide the basic insights about the final result of the interpretation. This paper is mainly concerned with the first of these three aspects.

2. OVERVIEW OF THE PHENOMENA TO MODEL

This paper is concerned with the interpretation of noun phrases and in particular with NP's occurring in intensional (or opaque) contexts and to the dichotomy existing between value-loaded and value-free interpretations. The final goal is to show how the formalism of tri-concepts (introduced in [Di Eugenio & Lesmo 87]), extended in a suitable way, accounts for the phenomena under study.

Particularly relevant in the present context seem the works by Fawcett and Hirst [1986] and Kronfeld [1986]. The theory outlined in [Fawcett & Hirst 1986] implies the existence of seven different readings of the sentence

(5) Nadia wants a dog like Ross's

Those readings depend on three different parameters:

- I) Transparent (de re) vs. opaque (de dicto)
- II) Source of description (the agent, i.e. Nadia, vs. the speaker)
- III) Time of application (present vs. future vs. repeated)

We believe that only three readings are directly related to the informational content of the sentence, since neither the source of the description, nor the time of application seem to be among the concerns of the sentence itself. For instance, Fawcett and Hirst distinguish between

a) Transparent reading, agent's description

(Nadia says "I want that dog, which is just like Ross's". The speaker reports this as (5) perhaps without being familiar with Ross's dog)

b) Transparent reading, speaker's description

(Nadia says "I want that dog", and the speaker, who knows Ross's dog, reports this as in (5)).

The problem is that (5) does not seem to be concerned with what Nadia or the speaker know about Ross's dog; it only assumes that "a dog like Ross's" is a description valid for the hearer of (5) so that it can enable him to build an internal representation of what Nadia wants. On the other hand (5) certainly involves a description of Nadia's desires; taking this into account, a perspicuous representation should be able to characterize those desires more precisely: Does Nadia wants a dog of a kind that casually happens to be the same kind of Ross's dog? Or the dog she wants must be of that kind just because it is the same kind of Ross's dog? In other words how relevant is the description?

The relevance of descriptions seems to be one of the central concerns of the work described in [Kronfeld 86]. In that paper, a computational model of reference is provided, and definite descriptions outside intensional contexts are considered. Some examples are:

⁵ "The construction of spaces should not be over-interpreted. They are not representations of reality or of partial 'possible worlds'. In particular, a belief space ('Max believes...') is not a belief or state of belief; at most, it is a way of talking about beliefs". [Fauconnier 85, p.152]

⁶ "The space elements are not referred to by expressions of language; they are set up, identified, etc., and may then be used for purposes of reference (or possibly purported reference). However, this distinction has an obvious and yet nontrivial consequence: theories of reference based on the use of language cannot bypass mental spaces; they will have to forsake the idea of a direct link between linguistic structures and referents and take into account the important intermediate process of space construction". [Fauconnier 85, p.158]

(6) Smith's murderer is insane

(7) The winner of the race will get \$10000

Let's consider ex. (6): traditionally, "Smith's murderer" can be used referentially in case the speaker "knows" who is the murderer; otherwise, it must be used attributively (i.e. when "whoever he is" is appropriate)⁷. Note that, even in case the speaker does not have any clue about the identity of the culprit, still the expression refers to a single, well defined individual. In fact, one of the accomplices of the murderer, after hearing a policeman say (6) on a TV interview, could report that as "Hey, Bob, he told that *you* are insane". Analogously, the fact that no winner of the race exists yet does not mean that the reference in (7) is not definite or precise, since the "role" of winner is uniquely defined even when the race has not been run. Again, if Fred wins, it is possible to report (7) telling him "*You* will get \$10000". Notice, however, that in both examples the reports are correct but not precise, since they fail to convey the "reason why" the thing reported is true. So it seems to us that in "attributive" readings there is a component of meaning concerning some kind of dependency (Causality?)⁸. It is worth noting that this proposal is completely compatible with Barwise and Perry's. In [Barwise & Perry 83], in case of attributive reading, the description must be evaluated in all the situations which contain an individual who matches such a description. That means that the assertion made about this description must be relevant in all these situations. This implicitly states that there is a sort of dependency connection between the description and the assertion, and that this connection is valid in all possible situations which contain an individual satisfying the description. What we aim to show is that the distinction between attributive and referential readings provides us with an important criterium that also affects intensional readings.

We believe that the two "opaque" readings we envisioned earlier are distinguished by means of the same "relevance" criterium used above. In classical examples as

(8) Mary wants a dog

no relevance criterium can be applied because of the simplicity of the description. In fact, if Mary said "I want a nice and faithful animal", either we assume that, for Mary, the intension of "nice and faithful animal" coincide with the intension of "dog" (i.e. in all possible worlds the two expressions denote the same set of individuals) or the report is not accurate. On the contrary, we could assume that "a dog like Ross's" is right only in this world (in any world she wants the same kind of dog, which is the same as the kind of dog that Ross owns in this world) or in every world (whichever is the kind of dog Ross owns, in any world, Nadia wants to have one of the same kind). Let's consider one more example to make the point clearer (see [Fauconnier 85, pp.45-46]):

- George says: "The winner will go to Hong Kong"
- The speaker assumes that "the winner" and "the best dancer" are equivalent descriptions (or "equal roles in the speaker's mental space" in Fauconnier's terms)
- The speaker says "George thinks the best dancer will go to Hong Kong"

We claim that the speaker's report is incorrect (although sincere), contrarily to what Fauconnier assumes, because in an "attitude" context (as "thinks" in the present example and "wants" in the previous case) what must be correctly modelled is the attitude of the agent. The speaker's description always has a subsidiary role: the one of describing correctly the mental state of the agent.

3. LAYOUT OF THE SYSTEM

This short section provides some information about the overall organization of the system which the semantic interpreter is a part of. The GULL (General Understander of Likely Languages) system works on the Italian language and is based on production rules. The syntactic structures are represented as dependency trees, where all nodes, and not only the leaves are associated with surface words. For

⁷ What "knowing" means in the present context is not very clear. If the speaker determined the identity of the murderer by finding his driving license near the corpse and this is the only thing he knows about him, then it can safely be stated that he "knows" who is the murderer, but still the only possible reading is the attributive one.

⁸ "Con conversationally relevant use" in Kronfeld's terms.

each subtree, the root represents the "head" of the constituent and its immediate subtrees are the "modifiers". The syntactic production rules say how the tree must be extended when a word of a given syntactic category is analyzed (for more details see [Lesmo & Torasso 83, 85a, 85b]).

As soon as the syntactic rules extend the tree, the control is passed over to the semantic rules: they update the semantic representation according to the meaning of the new piece of tree (in other words, syntax and semantics operate synchronously; and the knowledge sources are separate). Backtracking is not used: whenever the analysis is blocked, the current situation is inspected by another set of rules ("restructuring rules") that modify the dependency tree in order to provide a new interpretation; then, the analysis is restarted on the new tree. Semantically ambiguous words that belong to the same syntactic category are disambiguated on the basis of the local context where they appear: no possibility of undoing the choice on the basis of a larger context is currently provided. The antecedents of the semantic rules pick up any newly entered word (the "current node of the tree") and match it, together with its surrounding context (i.e. its mother node and its dependants), against the domain knowledge, represented in terms of a KL-ONE-like semantic net, which fulfills the role of the terminological component in an hybrid knowledge representation system. The access to the net binds some variables appearing in the antecedent of the rules to nodes of the net (concepts), which are then used in the consequent to build up the representation that is going to be described. Of course, the import of non-content words (as articles and quantifiers) is different from the one of content words. In a previous paper [Di Eugenio & Lesmo 87] we described how the introduction of non-atomic nodes called tri-concepts enabled us to deal deterministically with definite and indefinite determiners. Now, we will see how the problems about intensional contexts described before can be solved. In parallel, we worked on quantifiers: some results have been presented in [Lesmo et al. 88].

The semantic representation built by the interpreter consists in two planes (CP: Content Plane; RP: Reference Plane). A third plane is used in the drawings to refer to the stable domain knowledge (the terminological component mentioned above, called SP - Semantic Plane). The CP includes a set of nodes and arcs that represent the main assertion and the description of the entities involved in that assertion (i.e. the restrictions imposed on them by prepositional phrases, relative clauses, etc.). The RP contains a set of non-atomic structures called RASs (Reference Ambiguity Spaces) composed of a Concept node, an Extension node, and a Class node (three aspects of a "concept" that are needed to establish the possible referent of a linguistic expression). These structures will not be used often in the present paper, since they are more concerned with the representation of articles. A fourth node can be present in a RAS: it represents a subset of the extension and is used when a specific reading is chosen for definite noun phrases. On the contrary, since indefinite determiners are not used in class assertions, the Class node does not appear in RAS's associated with indefinite descriptions. Each RAS is linked to the defining concept of the SP via a SEMREF (SEMantic REFerence) arc. Also in the CP there may appear complex structures (DAS: Distributivity Ambiguity Spaces): They contain two nodes (Set and Individuals) and are used to distinguish between distributive and collective readings. The arc which connects a node in the CP to its associated RAS in the RP is labelled SAME for definite NPs (it simply states that the given node of the CP coincides with the reference node in the RP), I-OF (Instance-OF) for indefinite ones.

In order to simplify the figures, we will use the conventions shown in fig.1. Note that the initial representation (as shown in fig.1) will be augmented during the processing of the sentence if the NP has specifications (e.g. adjectives, relative clauses etc.). These specifications are not directly linked to the CP node that refers to the entity but to another node connected to the RAS by a DEF-AS (or NAME, if the specification is a proper noun) arc⁹.

⁹ The introduction of a pair of nodes for a unique referent is justified by the need of keeping apart the entity that participates in the main predication and the way it is defined.

4. REPRESENTATION OF INTENSIONAL CONTEXTS

The main extension to the work reported in [Di Eugenio & Lesmo 87] consists in the introduction of *CONTEXT SPACES (CS)*, which enable us to treat the intensional contexts along the lines proposed in [Fauconnier 85].

A Context Space is a subplane of the CP, which contains a set of nodes and of arc connecting them. CS are created when particular conditions occur in the input sentence, in our case the presence of an intensional verb. In this case the CS is created and connected to the node which represents the verb in the CP by a SP-BUILD (SPace-BUILDer) arc. The fact that the arc leaves a node representing an attitudinal verb specifies (implicitly) the type of the CS¹⁰ and its generator. The nodes appearing in a CS refer to entities that do not necessarily exist in the outer context (the CP itself is assumed to correspond to the outermost context, i.e. it contains terms that have to be interpreted according to the speaker's view of reality, as expressed in the SP). A node of a certain CS may refer to a local entity in that CS or to a counterpart in that CS of an entity of the outer context. The COUNP (COUNterPart) arc specifies the scoping of nodes in the semantic representation: an arc COUNP that connects a node N1 in a CS (say CS1) to another node (say N2) external to CS1 states that N2 is a counterpart of N1 (or, better, of the entity which N1 refers to) in the outer space; on the contrary, a COUNP arc starting from N1 and entering N1 itself specifies that N1 has no counterparts in outer contexts, i.e. it refers to a local entity of space CS1. The absence of a COUNP arc leaving from N1 represents an ambiguous case in which the scope of N1 has not been determined yet. Such ambiguous representation enable us, as we will later see in the paper, to develop a deterministic analysis of sentences also in case of de-re or de-dicto ambiguity.

Let's show how these structures are used by means of some of the previous examples. During the interpretation of sentence (5) the analysis of the fragment "Nadia wants" leads to the creation of a new CS (CS1 in fig.2). CS1 represents the "world of Nadia's desires as mentioned in the sentence" and is linked to node N2 (instance of "to want") through a SP-BUILD arc. Node N3 represents what Nadia wants as an instantiation of the "dog" concept and is introduced in CS1 when "a dog" is analyzed. Up to this point, the representation must be considered as ambiguous, since it does not specify which is the status of N3 with respect to the speaker's world. In case no further processing is possible, then the sentence remains ambiguous (fig.2 shows the basic ambiguous representation of sentence (5)); in other cases, perhaps depending on the context where the sentence is uttered, it could be possible to decide whether the intended reading is

¹⁰ In this paper we will only deal with "attitudinal" CS's. Actually many other types of CS's may be used to cope with other kinds of linguistic phenomena. For example a "temporal" CS may be used to deal with the ambiguity of "In 1970, the president met my sister", where "the president" could refer either to the current president or to the person who was president in 1970.




Fig.1 - Basic structures used in the representation

"de-re" or "de-dicto" (also preference criteria could be adopted; for instance, it seems that, outside any specific context, the preferred reading of (5) is the intensional one); when such a disambiguation is possible, then the representation is expanded by including a suitable COUNP arc: it will leave from node N3 and will enter node the same node N3 in case the reading is intensional, otherwise it will enter a newly-created node outside CS1 (the latter alternative is reported in fig.3). It is clear that it is the presence of node N8 that relates the representation to the existence (in the speaker's world) of a specific entity wanted by Nadia, contrarily to what happens in the intensional reading, where no node (referring to a dog) existing outside CS1 is linked to N3.

As stated above, two different "intensional" readings can be obtained for sentence (5) depending on the relevance of the description. The representation we obtained till now corresponds to the purely referential reading of (5) (actually, this use of "referential" is slightly inappropriate, but the previous discussion should enable the reader to interpret it correctly). If we want to represent the "attributive" reading we must add something to the structure; what we need is a specification of the "dependency" of the statement made through the utterance on the description used. To this aim we include the WHIS (WH-ever IS) arc, as shown in fig.4. At this point, the import of this arc should be clear.

The proposed formalism enables the system to operate deterministically, i.e. to build the final representation without destroying structures built previously (this was one of the goals of the introduction of tri-concepts, as explained in [Di Eugenio & Lesmo 87]). In fact the system builds a basic ambiguous representation (as shown in fig.2 for sentence (5)) which is possibly expanded with the introduction of new arcs when some disambiguation is possible. Of course, the insertion of suitable arcs does not modify the set of nodes initially stored in the CS's, which remain stable (perhaps expanded) throughout the analysis. As regards the disambiguation criteria, it does not seem that there are general rules for choosing one interpretation or the other. Of course, if the verb does not open an intensional context, then no problem arises (the CS is not created)¹¹; On the contrary, if we consider a verb as "seek", it seems that only the object NP can provide the clues to select one particular reading; in "Mary is seeking a street", the preferred reading certainly is non-intensional; in "Mary is seeking a friend" still seems to be preferable to assume the existence of a particular friend; on the contrary, in "Mary is seeking a funding" there seem to be few doubts that the funding does not exist yet. We are currently trying to build heuristic rules, based on nominal hierarchies appearing in the terminological component of an hybrid knowledge representation formalism, able to suggest a particular reading. Moreover, subsequent sentences may provide further disambiguation criteria. Consider the example in [Fauconnier 85], pag. 57:

(9) Ursula wants to marry a millionaire.

We may obtain different readings depending on the subsequent sentences:

(10) He is really a con-man

leads to a de-re (referential) interpretation, while

(11) But she won't find one

indicates a de-dicto reading for sentence (9).

A final point concerns the use of the proposed representation in the subsequent phase of interpretation. Generally, the search for the correct referent in the discourse history is led by the presence of a RAS in the RP and by its incoming (I-OF or SAME) and outgoing (a DEF-AS arc which connects it with the representation of the referent's description) arcs. An indefinite specific reference (for example "a dog" in "a dog is eating meat") causes the creation of a new entity, since the assumption is that indefinites are used when the referent is not identifiable by the reader. On the contrary definite specific reference (for example "the dog" in "the dog is eating meat") leads to the search for the entity referred to into the discourse history¹².

¹¹ Actually, the verb of the sentence is examined in all cases, since the presence of some auxiliaries forces the creation of an intensional context; consider, for instance: "You must reserve a seat", "I will build a table". Of course, these verbs, when used in the past tense do not have any intensional characterization.

¹² Obviously, many other types of references are possible. For a complete analysis of referential problem connected with the semantic of determiners see [Di Eugenio & Lesmo 87].

Let us see now the impact of the newly proposed extensions onto this environment. When a RAS is referred to by a node Nk through an I-OF arc (indefinite specific reference) the possible presence of Nk in a CS (say CS1) may be detected. If this is the case, node Nk must be inspected: if it has a COUNP arc, pointing to a node of the external context, then the usual creation process is executed; otherwise Nk refers to an entity which is local to CS1 and has no counterparts in the external world. So we need an interpretation which blocks any inference on the entity associated to the node in the external world. This result can be obtained by copying the whole CS in the final network representing the discourse history, and so hiding its internal node to the external context.

Let us specify now the interpretation of the WHIS arc. As stated above, when a RAS is referred to by a node Nk through a SAME arc (definite specific reference), an entity satisfying the description (represented by a Nj node of the CP connected to the RAS through a DEF-AS arc) is searched for in the discourse history

network. In fact, in this case, we only have to cope with "the functional relevancy of the description" [Kronfeld 86] (i.e. its ability to convey information leading to the identification of the correct referent). But what does it happen if we also have to account for conversational relevancy, i.e., what is the interpretation if a WHIS arc connecting the description node Nj to the node Nk exists? In this case the usual searching for the referent is delayed. Both node Nk and its description (represented by the WHIS arc and by the subnet of CP whose head is Nj) are entered in the discourse history network. In this way the description itself becomes part of the final representation of the sentence (as stated in [Barwise & Perry 83]). Moreover the presence of an arc WHIS in the final network states that the node entered must not be considered as the referent of the sentence but rather as a "role" (see [Fauconnier 85]) which may assume different values in different worlds (i.e. in different CS's). In fact, different referents for the description may be detected by the usual searching process in the CS's and this leads to different inferences (different assignments of referents to the "role").




Fig.2 -Basic ambiguous representation of the sentence "Nadia wants a dog like Ross's".

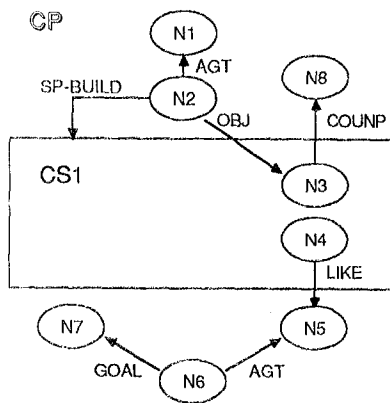


Fig.3 -Representation in the CP of the transparent reading of the sentence "Nadia wants a dog like Ross's"

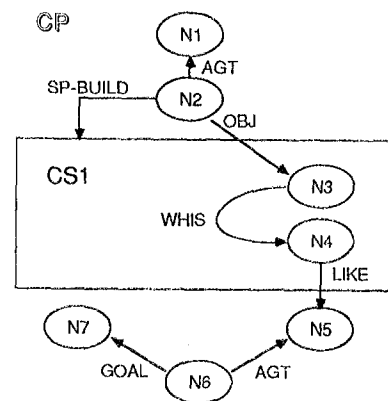


Fig.4 -Representation in the CP of the opaque attributive reading of the sentence "Nadia wants a dog like Ross's".

A final remark is needed: for the sake of simplicity we stated that the relevancy of the description must be detected during the analysis of the input sentence, independently from the subsequent interpretation process. Actually the strategy of the system is quite more flexible, because also the interpretation process, merging the semantic representation of the sentence and the discourse history network, supplies further information about the relevancy of a given description. For example, when the sentence does not include sufficient information to choose between referential and attributive reading), the absence in the discourse history of a referent fitting the description (detected during the final process of interpretation) forces an attributive reading.

5. CONCLUSIONS

We have described a semantic network formalism that is able to cope with the problems arising when NP's occur as objects of verbs inducing an intensional context. The representation preserves in an easily accessible format all information needed to distinguish among the possible readings, and the construction of the net can be carried on deterministically (see the discussion of paragraph 3 and note that disambiguating among different readings only consists in adding arcs to existing subnets). However, some problems still have to be solved concerning the final representation obtained by the process of merging the meaning of the sentence under analysis with the discourse history. Actually, the "assertional" component of hybrid formalisms (see, for instance, [Brachman et al. 83]) does not lend itself to handling intensional objects created during the analysis. As stated above, we have structured the assertional component providing facilities for copying the subnetworks containing the description of intensional entities into it. On the other hand, some work still has to be made in order to provide such a component with a clear semantics. We believe that basic insights can be found in the SNePS system [Maida & Shapiro 82], [Shapiro & Rapaport 86], where the proposed organization meets more clearly the requirements set forth by the existence of intensional objects.

The version of the interpreter described in this paper is implemented in FRANZ LISP and runs on a SUN workstation under the UNIX operating system. The maintenance of the discourse history and the merging process is currently under implementation.

ACKNOWLEDGMENTS

This work has partially been funded by the Ministero della Pubblica Istruzione of Italy, within the project "Software Architectures for Intelligent Systems". We must thank Kai von Luck and his colleagues of the Technische Universitaet Berlin, for having put at our disposal BACK, a hybrid Knowledge representation system whose terminological component has been used to implement the Semantic Plane. The second author (Paolo Terenziani) acknowledges Regione Piemonte, Progetto Finalizzato Sanita', that provided the funds for his Research Assistantship.

REFERENCES

- [Barwise & Perry 83] J.Barwise, J.Perry: *Situations and Attitudes*. MIT Press (1983).
- [Brachman et al. 83] R.J.Brachman, R.E.Fikes, H.J.Levesque: KRYPTON: Integrating Terminology and Assertion. *Proc. AAAI-83*, Washington, D.C. (1983) 31-35.
- [Cooper 83] R.Cooper: *Quantification and Syntactic Theory*. D.Reidel Publ. Co. (1983).
- [Di Eugenio & Lesmo 87] B.Di Eugenio, L.Lesmo: Representation and Interpretation of Determiners in Natural Language. *Proc. 10th IJCAI*, Milano (1987) 648-654.
- [Dowty et al. 81] D.R.Dowty, R.E.Wall, S.Peters: *Introduction to Montague Semantics*. D.Reidel Publ. Co. (1981).
- [Fauconnier 85] G.Fauconnier: *Mental Spaces*. MIT Press (1985).
- [Fawcett & Hirst 86] B.Fawcett, G.Hirst: The Detection and Representation of Ambiguities in Intension and Description. *Proc. 24th Annual Meeting of the ACL*, New York (1986) 192-199.
- [Kronfeld 86] A.Kronfeld: Donnellan's Distinction and a Computational Model of Reference. *Proc. 24th Annual Meeting of the ACL*, New York (1986) 186-191.
- [Lesmo & Torasso 83] L.Lesmo, P.Torasso: A Flexible Natural Language Parser Based on a Two-level Representation of Syntax. *Proc. 1st Meeting of European Chapter of ACL*, Pisa (1983) 114-121.
- [Lesmo & Torasso 85a] L.Lesmo, P.Torasso: Analysis of Conjunctions in a Rule-based Parser. *Proc. 23rd Annual Meeting of ACL*, Chicago (1985) 180-187.
- [Lesmo & Torasso 85b] L.Lesmo, P.Torasso: Weighted Interaction of Syntax and Semantics in Natural Language Analysis. *Proc. 9th IJCAI*, Los Angeles (1985) 772-778.
- [Lesmo et al. 88] L.Lesmo, M.Berti, P.Terenziani: A Network Formalism for Representing Natural Language Quantifiers. *Proc. ECAI 88*, Muenchen (August 1988).
- [Lewis 72] D.Lewis: General Semantics. In D.Davidson, G.Harman: *Semantics of Natural Language*. D.Reidel Publ. Co. (1972) 169-218.
- [Maida & Shapiro 82] A.S.Maida, S.C.Shapiro: Intensional Concepts in Propositional Semantic Networks. *Cognitive Science 6* (1982) 291-330.
- [Shapiro & Rapaport 86] S.C.Shapiro, W.J.Rapaport: SNePS Considered as a Fully Intensional Propositional Semantic Network. *Proc. AAAI-86*, Philadelphia (1986) 278-283.