NON-SINGULAR CONCEPTS IN NATURAL LANGUAGE DISCOURSE

Tomek Strzalkowski

Department of Computer Science
Courant Institute of Mathematical Sciences
New York University
251 Mercer St.
New York, NY 10012 U.S.A.

Nick Cercone

Centre for System Science
School of Computing Science
Simon Fraser University
Burnaby, British Columbia V5A 1S6 Canada

We introduce a new approach to representing and manipulating various types of non-singular concepts in natural language discourse. The representation we describe is based on a partially ordered structure of levels in which the objects of the same relative singularity are assigned to the same level. Our choice of the representation has been motivated by the following main concerns: 1. The representation should systematically distinguish between those language terms that are used to refer to objects of different singularity, that is, those classified within different but related levels of the model; 2. The representation should capture certain types of inter-sentential dependencies in discourse, most notably anaphoric-type cohesive links; 3. Finally, the representation should serve as a basis for defining a formal semantics of discourse paragraphs that would allow for capturing the exact truth conditions of sentences involving non-singular terms, and for computing interlevel inferences. In this paper we discuss (1) and (2) only. (3) is currently under investigation and will be the topic of a forthcoming article. We believe that our approach promotes computational feasibility, because we avoid the identification of general terms, like "temperature," "water," etc., with intensions, that is, functions over possible worlds. In our theory, the concept of non-singularity has a local (often subjective) character.

1 INTRODUCTION

Treating non-singular concepts is a difficult representation problem in natural language research. Non-singular concepts, as we shall understand them here, are often abstract entities created to embrace a variety of smaller or larger collections of "instances" or "specimen." They are usually referred to using bare plural noun phrases (such as birds, alligators, presidents), or definite singular noun phrases with "generic" interpretation (the alligator, the president). The literature describes numerous forms of non-singular concepts and corresponding to them non-singular terms that can be found in natural language discourse, including intensional (or functional), mass, generic, habitual, abstract, and more; see, for example, Montague (1974a), Lewis (1976), Barwise and Perry (1983), Quine (1960), Donnellan (1971), Vendler (1971), and Kripke (1972). In these and other writings, various treatments for the phenomenon are suggested, but many of them do not properly capture the distinction between singular and non-singular interpretation of linguistic descriptions. With the exception of intensional concepts, other forms of non-singularity have not been given satisfactory formal representations that would account for their role in natural language discourse. Perhaps the most successful treatment of non-singular terms in language thus far has been presented by Montague (1974a-d) with his formalization of intension, which can be traced back to the Fregean notion of sense. Unfortunately, the concept of intension does not capture all aspects of non-singu-

larity, for example, bare plurals denoting kinds cannot be adequately represented in Montague's intensional logic (IL) without some superficial, and often unplausible, extensions (Carlson 1977, 1982). In addition, the enormous complexity of any non-trivial system of possible worlds has proven to be disadvantageous for developing a computationally oriented application of Montague's theory (Strzalkowski and Cercone 1986).[1] In this paper we do not attempt to improve on Montague's grammar; in fact, our approach does, in part, rely on Montague's method of coupling the syntactic and semantic processing in order to assign to a sentence its meaning or at least a meaning-representing formula. Nonetheless, we also perceive the limitations of Montague's method. It is not our intention here to build an alternative formal system for a fragment of English that would replace that of Montague's. Instead, we propose a computationally oriented approach to certain problems with which Montague dealt only marginally or not at all. However informal our presentation may be, we make some effort to put things into a clear, semiformal setting. As various aspects of out theory crystalize, we expect a more formal version to emerge.

We introduce a fragment of a new and, we believe, computationally feasible theory of names and descriptions, which offers a uniform treatment for many types of non-singular concepts found in natural language discourse. Although we limit our presentation to nominal phrase constructions, the approach can be further extended to cover other types of phrases. We present the formal definition of non-singularity with respect to a particular discourse situation involving a discourse message; a number of individuals (parties); and their knowledge, beliefs, awareness, etc. We introduce a layered model of reality (the universe) as perceived by a discourse participant, and define relative singularity of objects in this universe as an abstraction class of the layer-membership relation. Subsequently, linguistic descriptions and names are classified as **singular, measurably singular**, or **non-singular** depending upon what they are assumed to denote in the universe. The relationship between objects referred to in discourse and classified into different layers (levels) of the universe model has a particular significance for resolution of certain types of cohesive links in text. We call these links **remote references** because they cross level boundaries. In order to adequately explain the phenomenon of a remote reference, we propose a fairly general theory of names and descriptions in discourse. Although the theory lacks a complete formal specification at this time, we introduce a number of notions, such as a **superobject** and a **coordinate**, that, we believe, will prove helpful in further formalization of various intuitions presented here.

In its present form, our theory focuses entirely on the problems of non-singularity and remote references in discourse that are created by the use of definite descriptions, bare plurals, and proper names. We do not describe in detail how one could actually compute

remote references in discourse, except for two general rules introduced in Section 6. These rules, and perhaps some others yet to be developed, create a natural expansion of the system for computing intersentential anaphoric references described in Strzalkowski and Cercone (1986). Except for a brief overview in Section 2, we do not discuss many related phenomena, such as non-remote (or single level) references, opaque readings, non-referential interpretations, use of personal pronouns, and the like. The more general theory of stratified meaning representation (Strzalkowski 1986) addresses all of these concerns, as well as the problems of discourse coherence, selecting proper cohesive links in discourse, and building a discourse model.

2 COMPUTING INTERSENTENTIAL DEPENDENCIES IN DISCOURSE

A first step toward automating the process of discourse understanding is to grasp the meaning contents of the discourse message, at least the literal meaning. A discourse normally consists of more than a single utterance, and although every utterance may be assumed to contribute something to the discourse meaning as a whole, this latter can only rarely be regarded as a simple sum of meanings of component utterances. Utterances, or sentences, making up a discourse are usually involved in complicated mutual dependencies, that often go beyond the text itself. A careful study of these extrasentential and intersentential dependencies in discourse is necessary before a more successful attempt to design an automated discourse understanding system can be undertaken.

In Strzalkowski (1986) and Strzalkowski and Cercone (1986) we introduced a rigorous method for handling certain cases of extrasentential dependencies in discourse, within a general framework, which we call the **Stratified Model**. The Stratified Model comprises a collection of disambiguating transformations that are applied to a discourse fragment before it can be assigned a final representation. These transformations include morphological analysis, lexical disambiguation, syntactic parsing, computing extrasentential dependencies, and pragmatic evaluation in discourse context. In order to compute anaphoric, and other, dependencies between sentences in a discourse fragment, we first translate each sentence in the fragment into an intermediate, formal representation language. We call these translations literal or "context-less" because they are derived based solely on sentences' syntactic structure. In fact, this transformation is done Montague-style (Montague 1974d) except for the intermediate representation, which, in the Stratified Model, is a λ-categorial language Λ, rather than an intensional logic. For this presentation it is enough to say that Λ is a typed predicate-calculus language with λ-operator. A formal definition of Λ is given in Strzalkowski and Cercone (1986). Next, we

proceed to compute various contextual dependencies of sentences, including extrasentential anaphoric links.

Let L be a language of parse structures obtained in some grammar of a fragment of English. For the purpose of this presentation, L is identified with the set of phrase markers that can be generated from English sentences with a categorial grammar CAT similar to that of Montague (1974). Although some other syntactic system may be more suitable in practical application, we select CAT here for its relative simplicity and elegance. We concentrate on the translation of some example expressions, sentences, and paragraphs of L into a representation in a λ-categorial language Λ that would capture both a sentence's logical form and its cohesive links to the surrounding discourse. In particular, we shall look closely at the cohesive links created by inter-sentential anaphoric references appearing in different contextual situations.

Λ possesses adequate expressive power to represent the meaning of a considerable spectrum of linguistic constructs found in a natural language discourse. What is of a particular interest to us, Λ provides a natural and uniform means for computing and representing extra-sentential dependencies. As we shall see in the next section, a meaning representation language so-defined is still inadequate for capturing some more difficult cases, which we call remote co-references.

Our present effort is to describe a transformation ISD such that ISD $\subseteq$ L $\times$ Λ, and whenever a source expression in L consists of more than one sentence, a class of intersentential dependencies within this fragment is identified and resolved, if possible. It must be noted here that ISD represents a semantic process that is entirely independent of any pragmatic or domain related factors. As a result a substantial amount of domain oriented ambiguity may be left unresolved. In any practical application, this transformation must be accompanied by a pragmatic process, as described in Strzalkowski (1986). ISD consists of a collection of translation rules $\{R_1, R_2, \ldots\}$, such that each rule is responsible for translating a specific type of dependency. Actually, only Rule 1 works directly on expressions of L, translating them into literal representations in Λ, independent of one another. Rules numbered 2 and up will take these literal translations and try to relate them pairwise looking, among other things, for unresolved anaphoric references. Most of these rules can be written in terms of two distinguished expressions of Λ, S_1 and S_2, which we call the **context-setting sentence** and the **current sentence**, respectively. Expression S_1 is a Λ-representation of the linguistic context in which the sentence with translation S_2 is to be evaluated. Neither S_1 nor S_2 must correspond to surface sentences, though. S_1 may represent a larger part of discourse, perhaps an entire paragraph; on the other hand, S_2 may constitute only a subclause of S_1 in which case we would talk of intra-sentential dependency. It should be noted here that the potentially explosive number of possibilities

will be in fact limited by the actual structure of the discourse under consideration (see, among others, Grosz and Sidner 1985), as well as by the pragmatic and domain related information, not discussed here.

Let us now consider a two sentence paragraph given below:

S_1: John interviewed a candidate.

S_2: The guy had impressive references.

In the most natural reading of this paragraph, the anaphor of "the guy" is resolved against "a candidate" in the first sentence, so that the second sentence actually means: "the guy whom John interviewed had impressive references." When considered separately from one another, S_1 and S_2 obtain the following translations into Λ.

$$S_1 \rightarrow \exists x[cand(x)\ \&\ int(J,x)]$$

$$S_2 \rightarrow \exists x[guy(x)\ \&\ C(x)\ \&\ \forall y[(guy(y)\ \&\ C(y))\supset$$
$$(x=y)]\ \&\ had\text{-}imp\text{-}ref(x)]$$

In the translation of S_2, C is a free predicate variable that needs to be bound by the sentence's context. The context variable is introduced into the translation of a definite noun phrase containing a definite article, such as "the," by Rule 1 (Strzalkowski 1986, Strzalkowski and Cercone 1986). This rule generates literal translations of sentences without considering their context, and does so in PTQ-style, that is, by assigning to each syntactic operation in CAT a formula formatting operation in Λ (Montague 1974a). The next step is to resolve context references; we must find a binding for C occurring in S_2 in the context provided by S_1 in order to obtain the final translation of the former. This intersentential dependency is captured by the translation Rule 2, which operates on the literal translations of both sentences delivered by Rule 1 (Strzalkowski and Cercone 1986). In the example above, the second sentence obtains the desired translation as shown below.

$$\exists x\ [guy(x)\ \&\ cand(x)\ \&\ int(J, x)\ \&\ had\text{-}imp\text{-}ref(x)$$
$$\&\ \forall y\ [\{guy(y)\ \&\ cand(y)\ \&\ int(J,y)\}\supset (x=y)]]$$

Rule 2 (Perfect-Context Translation Rule):
If the context-setting sentence S_1 has a referential interpretation in the form

$$\exists u[P(u)\ \&\ F(u)],$$

and the current sentence S_2 contains an unresolved definite anaphor, that is,

$$S_2 = \exists u\ [C(u)\ \&\ P_1(u)\ \&\ F_1(u)\ \&\ \forall x\ [\{P_1(x)\ \&\ C(x)\}\supset (x=u)]],$$

then this anaphor can be resolved against S_1, and the resulting translation of S_2 is obtained as

$$\lambda C[S_2](\lambda u[P(u)\ \&\ F(u)]).$$

A somewhat different problem arises when we consider a fragment with a possible non-referential interpretation, as in

John wants to marry a princess. The girl must be rich and pretty.

Now, Rule 2 can compute the anaphoric link between "the girl" and "a princess" only if both sentences receive their referential interpretations. In a case where both sentences are understood non-referentially, we have to use Rule 3, given below. No other combinations are possible.

Rule 3 (Imperfect-Context Translation Rule):

If the context-setting sentence S_1 has a non-referential interpretation in the form

imp $(\exists u \ [P(u) \ \& \ F(u)])$,

where **imp** *is an imperfect operator, and the current sentence S_2, also in a non-referential interpretation, contains a definite anaphor which occurs in scope of an imperfect operator* **imp**$_1$ *i.e.,*

$S_2 = $ **imp**$_1$ $(\exists u \ [C(u) \ \& \ P_1(u) \ \& \ F_1(u) \ \& \ \forall x \ [\{P_1 \ (x) \ \& \ C(x)\} \supset (x=u)]]$,

then this anaphor can be resolved against S_1, with the resulting translation of S_2 derived as

$\lambda C[S_2](\lambda u[P(u) \ \& \ F(u)])$.

Rule 3 encompasses a large class of non-referential contexts, which we call imperfect contexts, and which involve constructs including propositional attitudes (want, try, wish), intensional verbs (seek, conceive, think about), other complement-taking verbs (go, come), modal verbs (must, can, will), as well as progressive tense forms. In Rule 3, all this is reduced to the formula with the **imp** operator which translates compound phrases, such as "John wants," or "John will." Thus, in the example given above, Rule 3 is applicable when both sentences contain a wide scope **imp** operator. In this case, the second sentence of the fragment obtains the full translation with the following formula:

must$(\exists x \ [girl(x) \ \& \ princess(x) \ \& \ marries(J,x) \ \& \ rich(x) \ \& \ pretty(x) \ \& \ \forall y \ [\{girl(y) \ \& \ princess(y) \ \& \ marries(J,y)\} \supset (x=y)]])$

Other studied cases of intersentential anaphora (see Strzalkowski 1986a-c, Strzalkowski and Cercone 1986) include non-referential interpretation of discourse fragments involving attitude report verbs (believe, know, disagree). These cannot be translated with Rule 3, and a new rule, Rule 4, is developed to compute anaphoric links in texts similar to the one given below.

John believes that a unicorn lives in the park.
· He thinks the creature has a long horn.

Rules 5, 6, and 7 account for the pronominal anaphora, Rule 10 deals with certain instances of attributive use of definite noun phrases. Rules 8 and 9 are used when the antecedent of an anaphor is a proper name rather than a description. This is the situation where an interesting type of referential ambiguity occurs whose resolution

may have far reaching consequences on the process of discourse understanding.

Rule 9 (Names as Ultimate Referents):

If the context-setting sentence S_1 has the form of $F_1(N)$ where N is an individual constant denoting a name, and the current sentence S_2 contains a definite anaphor, so that its literal translation has the form

$S_2 = \exists x \ [P(x) \ \& \ C(x) \ \& \ F_2 \ (x) \ \& \ \forall y \ [\{P(y) \ \& \ C(y)\} \supset (x=y)]]$,

then the anaphor can be resolved against N as its ultimate referent with the following derivation:

$\lambda p[p(N)](\lambda x[\lambda C[S_2](\lambda s[\bar{N}(s)])])$

where $\bar{N}$ is the predicative use of name N.

In the following fragment,

Sylvester tries to catch a bird. The cat is clumsy.

there are two possible ways of linking "the cat" with "Sylvester." In one reading, not very different from those processed with Rule 2, the definite anaphor refers primarily to the entity that can be described as "the one who tries to catch a bird," and only contingently to its name. In this case we acquire some new information about Sylvester, namely that it is a cat. In the other possible reading, the anaphor refers to the name only, and thus may draw on some context that is different from the first sentence in the fragment. This latter situation is handled by Rule 9. In the above fragment, Rule 9 would produce the following translation for "the cat is clumsy" (S is an individual constant denoting the individual named Sylvester, and *Syl(x)* means that x's name is Sylvester):

cat(S) & Syl(S) & clumsy(S) & $\forall x$ [\{cat(x) & Syl(x)\} $\supset$ (x=S)]

There are more aspects of ISD transformation that merit attention. These include rules for dealing with other kinds of anaphora not discussed here, elliptical constructions, enumerably singular (plural) terms, intrasentential anaphora, and non-anaphoric dependencies, as well as indirect and forward reference cases where access to the speaker/hearer knowledge base may be required. We also have to deal with the changing reference level.

3 Non-Singular Terms in Discourse

The rules discussed in Section 2 cover selected cases of intersentential anaphora where the reference level in discourse does not change from one sentence to another. There exists, however, a class of intersentential dependencies whereby a reference is made across boundaries of different reference levels in discourse. For example, in

My new pet is an alligator. But the alligator cannot live in our climate.

"the alligator" in the second sentence most likely refers to a generic object of which the alligator in the first sentence is an instance or extension. Thus we can say that the second alligator is a non-singular superobject in which the first alligator somehow participates. The extent of such participation is not clear, but in general it can be observed that certain predications true of complexes of different kind are not preserved for their parts or elements, and vice versa. To represent this new kind of intersentential dependency we introduce a multilevel model for interpreting natural language expressions, such that the levels in the model would correspond (roughly) to the levels of reference in discourse. For instance, in the example above, the resulting representation would have both alligators placed at different, though related, "object levels." Because of an inherent subjectivity of such classifications, the levels in the model may have fuzzy boundaries and are only partially ordered with the "lower than" (i.e., "more detailed than") relation with respect to some current level (corresponding to the level of reference at a present point in discourse).

Consider now a somewhat larger fragment of text, excerpted from an article appearing in *The New Yorker* magazine.

> The two Ashanti kings are in somewhat different situations: the Ghanaian king is royally born, richly rewarded, divinely inspired and holds his office for life. The American Ashanti king is elected every two years from the ranks of an Ashanti social and cultural organization called the Asanteman Association of the United States of America, Inc. The first Stateside king, Kwadwo Tuffuor, was a plumber. The second, Kusi Appouh, repaired air-conditioners and refrigerators. Kwabena Oppong is the third king; he drives a cab.[2]

The first two sentences in this fragment contain direct references to the higher-level entities, which are the two Ashanti kings, as if they were ordinary singular objects. The remaining three sentences directly refer to the entities that are instances of one of these kings, now seen as superobjects, at different time intervals. The discourse reference level has changed, and now we talk about lower-level entities. The superobjects can still be referred to, but only indirectly, through their instances; this is what we call the remote reference.

The primitive notions of our theory are these of a singular object and a coordinate, a usually ordered set specifying a type of dimension that the object in question spans. A singular object is any entity to which we can directly refer using a nominal phrase of our language. The most common of the coordinates are time and space but other more abstract ones are also possible. These two basic notions are then used to define the notion of the object's *instance* with respect to some coordinate. Thus the pet alligator in the example above is related to the generic concept of alligator by some species coordinate that somehow ties (or enumerates?)

all alligators around the world. Similarly, if we use an appropriate coordinate consisting of two-year intervals we can decompose the American Ashanti king into its elected instances. If we reverse this process we can combine objects into complexes to which we can subsequently refer using collective terms, singular or plural, such as, for example, "people" or "the man" (generic). The *lower than* relation between levels in the universe model derives from expanding the notion of instance over collections of objects. The relation introduces a partial ordering within the universe model and thus helps to trace changes in the reference level of discourse. The highly discrete approach taken here is favorably contrasted with other existing approaches to non-singular terms, including Quine (1960), Kripke (1972), Montague (1974), Carlson (1982), and others. While insights of Quine, Kripke and, perhaps even more so, Carlson are undoubtedly of great influence, they require reworking in more discrete terms. Finally, we may note that the research in artificial intelligence and computational linguistics has devoted relatively little attention to treatment of non-singular terms in natural language in general and in natural language discourse in particular; see, however, Sidner (1979) for some early attempts to recognize generics in discourse. One of the goals of the present research is to fill this gap.

4 NON-SINGULAR TERMS IN LANGUAGE

Many philosophers and logicians, Quine (1960), Kripke (1972), Donnellan (1971), Vendler (1971), Montague (1974), and Barwise and Perry (1983), note that the usage of the italicized nominal phrases in Example 1 has a general or generic character, except for regular singular interpretations, which are only possible in some cases.

Example 1.
1a. *The king* wears a crown.
1b. *The president* is elected every four years.
1c. *Gold* is a yellow metal.
1d. *Temperature* is a measure of molecular motion.

Hundreds of similar examples involving such non-singular terms, and corresponding non-singular objects, as water, heat, the Pope, the number, etc. can be devised. Unfortunately, there is no generally accepted account of these non-singular terms in the philosophical literature. Some authors, for example, Vendler (1971) and Barwise and Perry (1983), cautiously called them generic, or general (for example, "the king"), or functional (such as "the number of students," "the temperature") uses of definite descriptions. Other authors, for example, Kripke (1972), were quite close to considering these kinds of non-singular terms as names (or at least some of them: heat, gold). Still other authors writing on the subject, for example, Quine (1960, 1973), advocated the notion of **abstract terms** as comprising of attributes, such as *(being)* red (further abstracted as "redness"), or

(being) the man drinking the martini (which cannot be so easily nominalized), which can be predicated of "concrete" objects. If we consider this discussion, the so-called "attributive use" of singular definite descriptions as identified by Donnellan (1971), may be considered as addressing some abstract, higher-level, and therefore (in our interpretation) non-singular concepts. Carlson (1982) discusses the case of so-called "natural kinds," a specific type among generic terms. He advocates the view in which generic terms are taken as denoting entities in the same way that singular terms do. If we accept Carlson's position, then our ontology of objects becomes far richer than before, and we need to impose more structure on our representation to reflect this new situation accurately. Indeed, Carlson introduces a special R relation into Montague's IL which allows him to create individual objects out of generic objects, as well as stages out of ordinary singular objects. Thus, "Max believes that dogs are here" receives the following translation, where *m* and *d* are individual constants denoting Max and dog kind, respectively (Carlson 1977):

$$\lambda x[Bel(\hat{\ }\exists y[R(y, x) \text{ \& } here'(y)](m)](d)$$

This formula says that Max believes that some stages of dog-kind are here. One problem with this representation is that we have no idea how these stages are to be identified, in other words, how do we decompose a kind-level object into stage-level individuals. If other types of non-singular terms denote in similar fashion, then we may need an even more complex model in which various stages (or levels) of object aggregation can be reflected.

Quine (1960) presents the most comprehensive account of various categories of terms found among natural language expressions. Almost everything that one can say is made up of different kinds of terms, appropriately connected to yield meaningful utterances, which he classifies as singular, general, relative, abstract, attributive, etc. At present, we do not draw such fine distinctions in our classification, reserving the right to develop extensions along the lines of Quine as needed. We present a few examples to provide additional insight and lay the foundation for our theory of names and descriptions.

There are numerous linguistic puzzles involving non-singular definite descriptions, among them Partee's (1972) famous temperature problem. Example 2 illustrates this phenomenon.

Example 2

2a. The temperature is rising.
    The temperature is ninety.
    Thus, Ninety is rising.[3]
2b. The president is elected every four years.
    The president is Reagan.
    Thus, Reagan is elected every four years.
2c. The tiger lives in the jungle.

    My pet is a tiger.
    Thus, My pet lives in the jungle.
2d. Americans drive big cars.
    John is an American.
    Thus, John drives big cars/a big car.

The arguments in (2a-d) are normally considered invalid. Various researchers agree that the definite descriptions "the temperature," "the president," and "the tiger" in the first sentences of (2a-c) should be interpreted functionally, that is, as intensions (Montague 1974d), or functions over situations (Barwise and Perry 1983). Note that if the descriptions were to be interpreted singularly or as enumerating all instances of a non-singular object (that is, statements containing them were understood as making claims about each instance), the reasoning would be valid. In (2d) the situation is somewhat more complicated, because every extension of the entity in the denotation of *Americans* is itself a compound entity, namely a generic entity. Thus, intensionality alone cannot explain why the first two sentences in (2d) appear connected. As a matter of fact, the reasoning displayed in (2d) is far more acceptable that those in (2a) to (2c). The reason, it seems, lies in our reluctance to apply singular interpretations to non-singular terms in first sentences of (2a-c), while this same move is more likely in (2d). In other words, while (2b') is an unlikely interpretation of (2b), (2d') can occasionally be accepted as a simplified version of (2d).

2b'. Every president is elected every four years.
    Reagan is the president.
    Thus, Reagan is elected every four years.
2d'. Every American drives big cars/a big car.
    John is an American.
    Thus, John drives big cars/a big car.

Of course, (2a) and (2b), but not (2c) or (2d), *could* be valid if the definite noun phrases in the second sentences (the temperature, the president) were understood as co-referential with appropriate noun phrases in the first sentences. This avenue is also closed, however, because we do not regard these phrases as co-referential, that is, they cannot be substituted for one another, nor can be their denotations. Compare this with (2e) below, where the co-reference between "carnivorous animals" and "these beasts" is easily made.

2e. Carnivorous animals live in these forests.
    These beasts are tigers.
    Thus, tigers live in these forests.

We claim here that no two descriptions can be considered co-referential unless they are used to refer to the objects that are at the comparable stages of aggregation. We call such object relatively singular and place them within the same level in our model. It is clear that the objects referred to by the corresponding descriptions in the first two sentences of (2a) to (2d) are not relatively singular, and thus the conclusions in the third sentences are not forthcoming. Another type of co-reference, which we call a **remote co-reference**, can still occur, and we put this view forward in this paper.

176

5 A MULTILEVEL MODEL FOR INTERPRETING NATURAL LANGUAGE TERMS

Initially, we note that our language tends to deal with singular objects only, no matter how complex their structure happens to be. A singular object is any entity that can be taken as a coherent whole, in other words, it can be referred to directly using a referring expression of language: a name, a definite description, a pronoun. Thus, at least as far as our ability to refer is concerned, all objects appear singular. Still, it is not the case that all objects are singular in the same way. Take, for example, two persons John and Mary. They are singular objects and they seem singular in the same way, in other words, singular relative to one another. Next take alligator, the species, and the alligator John owns. Although both are singular in their own right, they are not compatible when related to one another: the alligator John owns appears only a manifestation, or extension, of alligator the species at a certain space-time location. The individual alligator, which at some period of its life is owned by John is, therefore, singular in the same way John is, but this will not be the case when we consider the time slice of this alligator (an alligator stage, in Carlson's words) while it was owned by John. This latter appears only an instance of the individual alligator at some time interval. By the same token, if John owns many different alligators at different times (but, arguably, never more than one at a time) then we may risk to refer to John's alligator, an abstract object that generalizes over all alligator stages of all individual alligators ever owned by John (for example, "John always walks his alligator in the morning"). The new object, again, seems to belong in the same class of objects as John and Mary.

Let us introduce, only intuitively at first, the relation of **relative singularity** among objects. As suggested above, this relation will help us to break down the universe of objects into classes of relatively singular objects, which we call levels. The levels can be subsequently partially ordered with *lower than* relation, i.e., $L_1 < L_2$, indicating that level L_1 consists of manifestations (extensions, instances) of objects at level L_2. Let L_0 be an arbitrary level we select as our reference point; if our discourse operates at this level then L_0 defines the current level of reference of the discourse. Let L_{+1} and L_{-1} be two other levels different than L_0 and such that $L_{-1} < L_0 < L_{+1}$. At level L_{+1} we place the objects we consider to be generalizations (or abstractions) of some measurable amount of objects from L_0. It is only from the perspective of L_{+1} that we are able interpret "The tiger lives in the jungle," or "The president is elected every four years," or "Birds can fly," or "Tourists start forest fires." The objects at L_{+1} are singular but only when related to one another within the same level; when viewed from L_0 they appear generic or functional or the like, in other words, non-singular. When we attempt to find a denotation for a nominal, such as "the tiger" or "tourists," within L_0, we attempt to give it

either a singular or measurably singular interpretation. In a singular interpretation (if possible at all), we have a nominal refer to a specific object within the level (John's pet tiger), while in a measurably singular interpretation we use a quantification over a finite set of singular objects, also within the same level (every tiger, some tourists). Nominals denoting objects which are non-singular with respect to L_0, on the other hand, may have neither singular nor measurably singular interpretations within this level. In order to find proper denotations for them we must change the reference level from L_0 to L_{+1}. Thus, while the statement of "The President lives in the White House" when interpreted at level L_{+1} can be argued to be equivalent to the statement "Every president lives in the White House" interpreted at L_0, the same cannot be said of "The tiger lives in the jungle" and "Every tiger lives in the jungle." We must note that some objects found at L_{+1} could have been placed there by design rather than as a result of generalizing from L_0; an example of such higher-level object may be **The President**.

If level L_{+1} contains generalizations of objects from L_0, then level L_{-1} will contain their specializations or extensions. Descending upon L_{-1} we can see that what we previously considered to be the atom actually denotes many different kinds of atoms (H, O, Ca, Fe, etc.), or that the mail is not the same every morning, or that Nicolas Bourbaki is the name of a group of mathematicians.[4]

A few definitions will help to put the above intuitions into a more formal setting.

Def. 1. *A use of a description is called singular if it refers to a singular object. A use of a description will be called measurably singular if it refers to some measurable quantity of a singular object. Otherwise we shall talk of non-singular use.*

Def. 2. *An object level, or simply a level, is an arbitrary collection of relatively singular objects. On the language side, the corresponding reference level encompasses those singular and measurably singular uses of descriptions that refer to the level's objects.*

Def. 3. *For any level L, there are at least two distinct levels L_{-1} and L_{+1} such that L_{+1} contains these objects which are non-singular from the perspective of L, and L_{-1} contains the objects for which the objects at L are non-singular.*

Def. 4. *The level L_0 is an arbitrarily chosen level serving as a reference point.*

As described, the structure of levels is not yet adequate to capture the full complexity of the reference structure of discourse. A notion of coordinate has to be introduced along the following lines. We shall call T a **coordinate**, if T is a set of points or locations at which certain general (or abstract) objects, for example the president or the atom, are assigned more specific extensions or instances, such as President Reagan or H, Fe, Ca, ... A coordinate is usually an ordered set though

the ordering may be partial only. Almost any object we can think of appears an instance of a more general concept, and often there will be more such concepts available, if we consider different coordinates. Water in a glass is an instance of some totality of water in the universe (space coordinate), and also an instance of a concept of water as in "Water boils at 100 degrees Celsius." These examples suggest that a coordinate is usually a large set, often an infinite set, though perhaps no more than recursively enumerable. A non-singular object can be decomposed into instances in more than one way, depending which coordinate is used. An important observation is that when a higher-level object is decomposed with two different coordinates the resulting sets of instances need not belong to the same level. The concept of the American president in the twentieth century is an instance of The President, but is still non-singular when compared with President Reagan, also an instance of The President, though with respect to a different coordinate. A somewhat finer structure of levels is required.

Let $L_{-1}^{N,T}$ be the level where we place the instances of object N decomposed with coordinate T. By analogy, we define $L_{+1}^{N,T}$ to be the level such that for any object M, $M \in L_{+1}^{N,T}$ if $N \in L_{-1}^{M,T}$. In other words, $L_{+1}^{N,T}$ contains the superobject M generalizing over object N with the use of coordinate T. Suppose that we have an object N at level L_0, to which we refer using a description N.[5] Suppose further that coordinate T is selected so that for any $x, y \in T$ we have that N-at-$x \neq$ N-at-y. Let us use N_x to stand for N-at-x, where x is an element of T, and let (N x) be an expression (as translated into our meaning representation language) that refers to the object N_x, whenever the expression N refers to N. We obtain therefore that

F1. $\forall x,y \in T \, [x \neq y \supset (N \, x) \neq (Ny)]$

The new objects N_x's cannot be placed at L_0 because, being instances of N, they are not singular relative to N (see Def. 2). Instead, we move them onto a new level $L_{-1}^{N,T}$ leaving the original object N at L_0. We say that the level $L_{-1}^{N,T}$ is lower than the level L_0, and write $L_{-1}^{N,T} < L_0$. Often we drop the superscripts N and T over the level symbol, assuming some lower level L_{-1}, whenever it does not lead to ambiguity. Example 3 helps to illustrate the phenomenon just discussed.

Example 3. Let us consider a rather naive concept of bird, as that of a winged creature that lay eggs and can fly. Let B be extension of this concept in our model and let L_0 be set so that $B \in L_0$. Using a *genus* coordinate, G, we can construct a level $L_{-1}^{B,G}$ containing such objects as **eagle, hawk,** and **goose.** Let's suppose that, initially, **penguin** can also be found at level $L_{-1}^{B,G}$. Upon discovery that penguins cannot fly, however, one would wish to relax the characteristics of the concept B from L_0 to contain both flying and non-flying birds. Nonetheless, two distinct concepts emerge, that of flying birds and that of non-flying birds, both of which become subordi-

nates of the (now) more general concept B. These new concepts are placed at a new level $L_{-1}^{B,K}$, where $K=\{k_1, k_2\}$ is a class coordinate such that B_{k_1} is the concept of flying bird, **FB,** and B_{k_2} is the concept of non-flying bird, **NFB.** The old level $L_{-1}^{B,G}$ remains intact, though it is different from $L_{-1}^{B,K}$. Moreover, $L_{-1}^{B,G} < L_{-1}^{B,K}$ because $L_{-1}^{B',G_{B'}} \subseteq L_{-1}^{B,G}$, where B' is either **FB** or **NFB,** and $G_{B'} \subseteq G$. There is another way of interpreting concept B as well: we introduce a *specimen* coordinate S that allows us to pick up specific birds, such as Opus, the penguin, at level $L_{-1}^{B,S}$. Note that this level is lower than $L_{-1}^{B,G}$ because it contains all levels L_{-1}^{X,S_X}, where X ranges over objects at $L_{-1}^{B,G}$, and $S_X \subseteq S$. Note that the structure of levels has been created as described because of the following set of conditions.

$$L_{-1}^{FB,G_{FB}} < L_{-1}^{B,K} < L_0$$

$$L_{-1}^{NFB,G_{NFB}} < L_{-1}^{B,K} < L_0$$

$$L_{-1}^{FB,G_{FB}} \cup L_{-1}^{NFB,G_{NFB}} = L_{-1}^{B,G} < L_0$$

$$\text{Opus} \in L_{-1}^{penguin,S_{penguin}} \subseteq L_{-1}^{B,S} \text{ with } S_{penguin} \subseteq S$$

Figure 1 further illustrates the concept of levels and coordinates.[6] For an easy interpretation of this drawing, note that $g_1, \ldots, g_5 \in G$, $G_{FB} \cup G_{NFB} = G$ and $g_1, g_2, g_3 \in G_{FB}$ and $g_4, g_5 \in G_{NFB}$, and $s_1 \in S_{penguin} \subseteq S$.
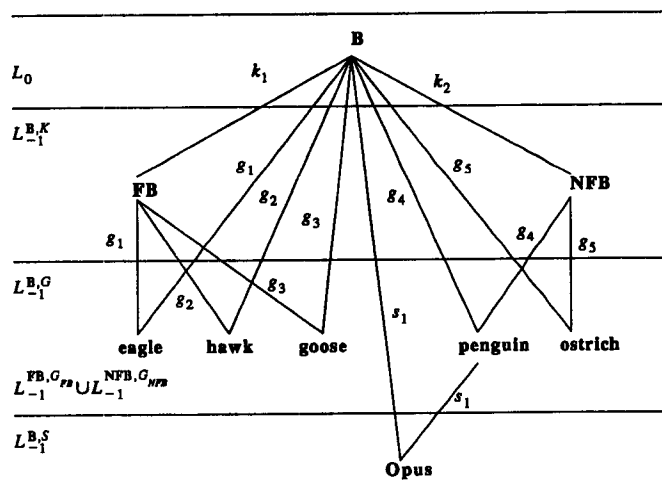

Figure 1. A structure of levels for a simple concept of bird.

Now we can attempt to represent meanings of some simple statements about birds. For example, *Birds can fly* is represented at L_0 as can-fly(B), while *Opus is a bird* would translate as $\exists s \in S \, [(B \, s) = \text{Opus}]$. We cannot infer from these statements that *Opus can fly* because this would require the formula $\forall s \, [\text{can-fly}((B \, s))]$ to be true, which does not have to be the case considering that B may now contain instances of non-flying birds. Indeed, *Opus cannot fly*, which translates to ¬can-fly(Opus), is not necessarily inconsistent with the above two.[7] The only thing we could say about each

instance of **B** is, perhaps, that it is a bird and that every bird is an instance of **B**. In other words, if x is a $L_{-1}^{B,S}$ variable then $\forall x$ [bird(x) $\equiv \exists s \in S[x=(B\ s)]$] is true. This would lead to an equivalent L_{-1} translation of *Opus is a bird* as bird(Opus). Later, we will see that this equivalence is not generally valid.

Let us now resume our general discussion on the characteristics of the structure of levels. At level L_{-1}, we have an enumerable collection of different objects N_x's. Extending the description used for N (at L_0) over N_x's we refer to them as *the N, a N, some N(s), every N*, etc. It is possible, of course, that some other object N' found at L_0 is now disclosed to be N_x, for some $x \in T$. What that means is that we have placed N' incorrectly at L_0, because it actually belonged to L_{-1}. Consider the case of an ancient astronomer who believed that the Morning Star and the Evening Star were not only two different heavenly bodies, but also of the same genre as other stars. In our conventions, both the Morning Star and the Evening Star were placed at the level to which the planet Venus now belongs. When correcting this ancient misconception, we faced the problem of an instance of some object and the object itself were mistakenly assigned the same singularity level; that is, we had both N_x and N at L_0. Nevertheless, this situation represented the state of our knowledge of the world at the time.

We may now give names to some N_x's and N may very well happen to be among them. This time, however, N will not denote the old object from L_0; this will be quite a different name referring the selected N_x and can be replaced by a definite description (N x). To illustrate this phenomenon, we may compare the concept of a programming language, such as Lisp, with its various implementations or versions available at different sites or times, and which are locally called Lisps as well. A more common transition of a name is, however, from an instantiation to a concept, which, ultimately, can create a similar effect. Consider for the moment the concept of a sun as a center (or one of several centers) of a solar system, and our own Sun as a specific instantiation of this concept. A schematic illustration of the descent process is shown in Figure 2.

. A process reverse to decomposition is that of ascending to a higher level within the level hierarchy. Suppose that for some objects N $_1$, N$_2$, $\cdots$, considered distinct at L_0, we discover they share a certain property, such as being an N, so that we need a generalizing concept to talk about them. We pick up a coordinate T, and climb onto some higher level L_1, that is, $L_0 = L_{-1}^{N,T} < L_{+1}^{N,T} = L_1$, and establish a new object N there, a "superobject." Now, as viewed from L_1, all N_i's are just the occurrences of N at different values of coordinate T. In other words, the following equation holds:

F2. $\forall i\ [\exists x \in T[(N\ x) = N_i]]$

Note that all N_i's now belong to the level $L_{-1}^{N,T}$, which is a part of L_0. As before, we shall drop superscripts N and
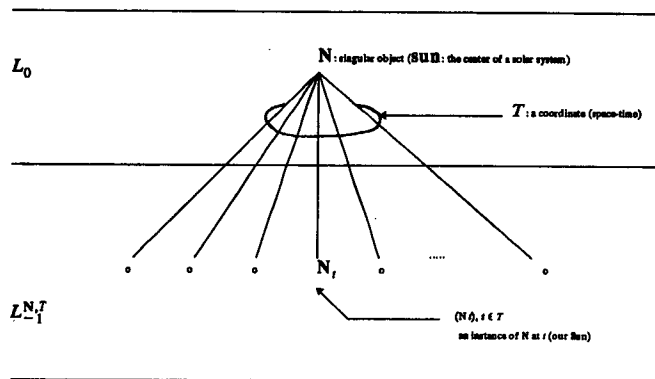

Figure 2. Descent from the current level onto a lower level.

T for simplicity. No matter how we name N at L_{+1}, the following **Formula of Discovery** summarizes our action:

F3. $\forall x,y \in T\ [(N\ x) = (N\ y)]$

It must be noted that (F3) is valid only when stated from the perspective of L_{+1}. At L_0, N_i's remain distinct traditionally, so they remain distinguished in the language as well; cf. Formula F1. The generalization of other objects from L_0 onto L_{+1} may follow but, as in the case of decomposition discussed earlier, the process will largely remain implicit. These observations are illustrated in Example 4.

Example 4. At level L_0, we have object **TP**, named The President. Let T be the time coordinate (which is different than T in the last two examples). At L_0, we have, according to (F3), that $\forall x,y \in T\ [(TP\ x) = (TP\ y)]$. Later, we may discover that for some t_1, $t_2 \in T$, (TP t_1)=N and (TP t_2)=R, and that at some level $L_{-1}^{TP,T}$ where N and R belong, they are considered distinct and named Nixon and Reagan, respectively. But at L_0, R=N is true. The last observation can be made clearer if one imagines that **TP** is some abstract individual who, when observed in the early 1970s, is named Nixon, and who, when observed in the 1980s, is named Reagan.

As described, the case of generalization (or abstraction) is extremely common in natural language, and the best tangible manifestation of this phenomenon is common nouns. Common nouns should be considered names of generalized concepts, which may be either of a physical or otherwise measurable nature (mass concepts such as water, snow, gold, temperature, . . .) or of a non-physical, abstract nature (usually based on enumerable quantities of specific instances: tree, man, president, . . .), or both (like fish, people, . . .).

Once the superobject N has been created, it begins to live a life of its own. Some new objects from L_0, different than N_i's, may now become instances of N at some, as yet, unutilized values of coordinate T. Also, we may use descriptions (N x) without caring whether they actually refer to any objects at L_0. In other words, the fact that a superobject N has no instance at certain location $x \in T$ at level $L_{-1}^{N,T}$ does not preclude the use of

the description (N x) in the language. The resulting expression may not always be well defined, though, as is in the case of "the present king of France," where the adjective "present" specifies the value of time coordinate decomposing the superobject **the king of France**. This problem of non-denoting expressions created out of denoting general terms, which is widely discussed by Quine (1960, 1973), receives an elegant explanation in our theory. One remaining problem is the relationship between a superobject and its instances in some given decomposition. It is important that we do not equate a superobject at L_{+1} with the set S of its instances at L_0, even if they may have actually given birth to this superobject. A superobject cannot be understood as a set of appropriate lower-level instances, since we would obtain only a measurable collection of singular objects. Instead, a superobject N can be identified with a family of functions $\{\Phi_T^N \mid T$ is a coordinate$\}$ such that each Φ_T^N is a function from coordinate T into an appropriate lower level, $L_{-1}^{N,T}$. In particular, a superobject N at $L_{+1}^{P,T}$, where $P \in L_0$, can be considered (from L_0 perspective) as a function Φ_T^N from T into L_0 such that, whenever $s \in S \subseteq L_0$, then there is $t \in T$ such that $\Phi_T^N(t) = N_t = s$. The function Φ_T^N is then arbitrarily extended beyond the set S. The following definition may be suggested.

Def. 5. *Let L and M be any two distinct levels of relatively singular objects. We say that level L is lower than level M, $L < M$, if there exists an object* **P** *at level M and a coordinate T such that $L \supseteq M_{-1}^{P,T}$.*

In order to avoid any misunderstanding, we add the following remark. It is possible that some N_j from among the N_i's was already recognized properly at L_0 as our goal object N from L_{+1}, although its other occurrences N_i for $i \neq j$ were not identified with it (cf. the Morning Star, the Evening Star, and Venus). In some sense, therefore, the previous concept of N was incomplete, since it did not contain these other instances which, in turn, allowed this former concept to coexist with some of its would-be instances at the same level. This fact can be further reinforced in our language when we choose to name N after N $_j$. This should not suggest that N and N $_j$ are one and the same object. The former is in a sense more mature, although, when referenced by name, one can hardly tell which one of the two concepts is being referred to, unless, of course, some additional clarifying context is present. Examples 5 and 6 below further illustrate this point.

Example 5 We have the following distinct objects at level L_0: V, called Venus; MS, called the Morning Star; and ES, called the Evening Star. Upon discovery that they all represent occurrences of the same planet, we create a new object V', named Venus, at the level $L_1 = L_{+1}^{V,T}$, and such that for some $x,y,z \in T$, where T is a time coordinate, $(V'x) = V$, $(V'y) = MS$, $(V'z) = ES$, where V', V, MS and ES are individual constants denoting V', V, MS and ES, respectively. Using the formula (F3), we conclude, from the perspective of L_1, that V = MS = ES, while the same conclusion made at L_0 is false.

Example 6. Let the level L_0 be as in Example 5, except that the object V is discovered not to be uniform. In fact, it contains occurrences of three different objects: planet Venus and some two heavenly bodies assumed to be Venus in the mornings and the evenings. Now we cannot use our time coordinate T from the previous example to get the desired result of the object V' at L_{+1}. Instead, we first descend to $L_{-1}^{V,S}$ over a coordinate S to differentiate the objects $V_{s_1}, V_{s_2}, V_{s_3}$ for some $s_1, s_2, s_3 \in S$. Let the V_{s_1} be a part of the ultimate object V'. In the same way, we create instances of **MS** and **ES** over the coordinate S at levels $L_{-1}^{MS,S}$ and $L_{-1}^{ES,S}$, respectively. Because **V**, **MS** and **ES** are relatively singular, their instances in decomposition with respect to the same coordinate yield objects that are also relatively singular. In other words, there is a level L_2 such that $L_2 \supseteq L_{-1}^{V,S} \cup L_{-1}^{MS,S} \cup L_{-1}^{ES,S}$. Let MS_{s_2} be called the Morning Star and ES_{s_3} be called the Evening Star at L_2. Now we can construct the ultimate object V' out of V_{s_1}, MS_{s_2}, and ES_{s_3} using a coordinate T such that $\{s_1,s_2,s_3\} \subseteq T$ and $V'_{s_1} = V_{s_1}$, $V'_{s_2} = MS_{s_2}$, and $V'_{s_3} = ES_{s_3}$. We place V' at a new level $L_3 = L_{+1}^{V_{s_1},T}$. Note that $L_3 \subseteq L_0$, and thus that V' is singular at L_0.

Example 6 is more "realistic" than Example 5, but both examples have equal linguistic significance. Note, however, that having L_2 as L_0, and L_3 as L_1, we could reconstruct Example 5.

6 REMOTE REFERENCES, SUPERCONTEXTS, AND SUBCONTEXTS

Let us examine how the foregoing theory of non-singular terms could be utilized in assigning meaning representation to natural language discourse. In particular, we are interested in the problem of computing extrasentential dependencies in text.

The translation rules described briefly in Section 2 are applicable only when anaphoric references arise between relatively singular descriptions or, as we would say now, between descriptions denoting objects classified within the same single level. We show that a similar technique can be used to compute another type of extrasentential dependencies, which we call the remote co-reference. Let us start with an example. The arrow symbol $\rightarrow$ is used here, and elsewhere, to mean "translates to".

Example 7. Consider the following discourse fragment.

    7a. The president$_1$ is elected every four years.

    7b. The president$_2$ is Reagan.

There is nothing to prevent us from interpreting the president$_1$ and the president$_2$ at levels L_1 and L_2, respectively, so that one of the following takes place. Either $L_1 = L_2$, or $L_1 < L_2$, or $L_2 < L_1$, or simply $L_1 \neq L_2$, where $<$ stands for the *extended lower-level* relation introduced in definition 5. The latter case does not interest us, since, in such an interpretation, both sentences were uttered at different occasions with no connection between them.

Consider first that $L_1 = L_2 = L_0$. If the two definite descriptions were to co-refer then we would be talking of the same object (individual) in both sentences. That interpretation, although possible, does not agree with our intuition. In this case the conclusion of

7c. Reagan is elected every four years.

follows immediately.

Let us assume next that $L_2 \supseteq L_{-1}^{TP,T} < L_1 = L_0$, where **TP** is the object at L_1 referred to by the president$_1$, and T is a time-based coordinate. If the president$_1$ is used as a name, we can expect to obtain the following translations:

7a $\rightarrow$ eefy(TP)

7b $\rightarrow$ $\exists t$ [SL(t) & ((TP t)=R)]

with "elected every four years" $\rightarrow$ eefy, "the president" $\rightarrow$ $\lambda P[P(TP)]$

where $t \in T$ and SL is a selector over T provided by the discourse situation (for example *now, here,* etc.). This selector could be obtained if we consider a possible translation of (7b) at L_2 as pres(R) & loc(R, now), from which it follows that $\exists t \in T[((TP \ t)=R \ \& \ loc((TP \ t),$ now)], with SL = $\lambda t[loc((TP \ t),$ now)].

In a more general case, we would take the phrase the president$_1$ as an ordinary definite description. Assuming some external context C which allows for the use of the definite article, the following translation can be derived:[8]

7a $\rightarrow$ $\exists r$ [p(r) & C(r) & $\forall y$ [(p(y) & C(y)) $\supset$ (r=y)] & eefy (r)]

7b $\rightarrow$ $\exists x$ [p(x) & C(x) & $\forall y$ [(p(y) & C(y)) $\supset$ (x=y)] & $\exists t$[SL(t) & p((x t)) & ((x t)=R)]]

where "president" $\rightarrow$ p.

The first of these formulas can be interpreted, with respect to some implicit context C (which may be reasonably assumed to limit our attention to the U.S. government), as: there is exactly one object such that it is a president and has a property of being elected every four years. Clearly, the object we refer to in (7a) is not any particular person, but the office of the President of the U.S. The other formula says, in turn, that there is a value t of the time coordinate T at which the instance of the general object referred to in (7a) (the President of the U.S.) is identical with the object **R** (which stands for Reagan). Observe that R refers to the individual Reagan as he appears at t, but not necessarily beyond that. The latter formula gives the full translation of (7b) when the remote reference to level L_1 has been resolved. Before this interlevel connection is established, however, we can assign (7b) only the a literal translation at L_2, as shown below, where C_2 is a local context at level L_2.

7b $\rightarrow$ $\exists x$ [p(x) & C_2(x) & $\forall y$ [(p(y) & C_2(y)) $\supset$ (x=y)] & (x=R)]

A point that is worth discussion is the role of the selector SL in the full translation of (7b). This selector may be interpreted as a local context at L_2 determined

by a discourse situation, as in Barwise and Perry (1983). To understand its influence on the form of the reference-making sentence, compare Example 7 with Example 8 below.

Example 8

8a. The tiger lives in the jungle.

8b. My pet is a tiger.

Unlike (8b), in the "story" of the president the local context allows (and requires) us to use "the" in (7b) because there may be at most one instance of the general term at a given discourse situation (in this case, the President of the U.S.). This is not the case when more than one value of the coordinate T satisfied the selector condition. The use of a definite description in such a context suggests, therefore, that a single instance of a general term is being picked up by the selector. This means, in turn, that except for a main "remote reference" in cross-level referencing, another local reference is being made at these levels where the definite description is used to denote some object. The remote reference itself does not need a definite description to be used for establishing a connection between an instance and the general term. The fact that we use a definite article when referring to an instance of a general concept, as in (7b), implies only that the local context SL contains (or is expected to contain) exactly one instance of the general object. Notice that the temporal aspect of this local context is extremely influential.[9] When we decompose a general object with respect to a time coordinate, we are more likely to obtain a unique instantiation in a local context. However, when a coordinate does not contain a time element, as in (8a,b), we cannot, in general, exclude the possibility that instances other than the one we intend to refer to in subsequent statements may be present in a local context. Thus a definite description is not used until the local context is properly narrowed. In the case of (8a,b) above, for example, we may continue the discourse at the lower level, saying "The animal is quite friendly," and meaning "The animal which is my pet tiger is quite friendly."

This discussion strongly supports our theory of names and descriptions, particularly the existence of levels and coordinates. Having established a higher-level object (or superobject), we can freely discuss its instances across various coordinates. The local references can be accounted for by singular translation rules, which we discuss in Strzalkowski and Cercone (1986).

Finally, let us define the notion of **remote reference** for objects classified into different naming levels.

> **Def. 6.** *An object* N *at a level* L_n *is said to be remotely referenced by a description* M *if* M *refers to an object* M *at some level* L_m *such that either* $L_m \supseteq L_{-1}^{N,T}$, *or* $L_m \supseteq L_{+1}^{N,T}$, *for some coordinate T.*

Let us summarize this discussion briefly now. In some part of a discourse, a certain (general) object **X** is addressed; that is, there is some part, S_1, of the discourse (presented as a single sentence in our examples,

for simplicity), such that S_1 predicates something of X—that is $S_1(X)$, where X is a description that refers to X. In a subsequent part of the discourse, however, the discourse changes the level of reference and only some instance(s) of X with respect to some coordinate T is addressed; that is, there is some $t \in T$ such that $S_2((X$ $t))$, where S_2 is this new part of the discourse. Apparently, the discourse internal cohesion would be compromised if we did not allow the higher level object X be a target of a remote reference by a description (X t) denoting one of its instances. In such a case we say that $S_1(X)$ creates a **supercontext** for (X t). We can further say that X and (X t) are **remotely co-referential**. By analogy, if S_1 contains a reference to an instance of an object X, that is, we have $S_1((X\ t))$ for some coordinate T where $t \in T$, and S_2 contains a reference to X, that is, $S_2(X)$, then $S_1((X\ t))$ is a **subcontext** for X. In this case too, X and (X t) are remotely co-referential.

We have just arrived at a very important property of a natural language discourse, which influences both coherence and cohesion, and is absolutely essential in proper representation of discourse meaning content. The following translation rule specifies how such a representation should be derived.[10] The approach presented here is entirely original, although the problem of the "general term/an instance" connection has been known for a long time and other, less general, solutions have been suggested (see Partee 1972 and Montague 1974).

Rule 11 (Supercontext Translation Rule):
If the context-setting sentence S_1 with the translation $\exists x\ [P_1(x)\ \&\ F_1(x)]$ is interpreted at level $L_{+1}^{\xi,T}$, where ξ is an object satisfying sentence S_2 when interpreted at level L_0, and S_2 contains an unresolved remote reference P_2, that is,

$$S_2 = \exists y\ [P_2(y)\ \&\ F_2(y)],$$

then the full translation of S_2 is obtained as

$$\lambda Q[\lambda C[M_{Q,C}](C_1)](\lambda u[\exists t[\lambda y[P_2(y)\ \&\ F_2(y)]((u\ t))]]),$$

where the supercontext C_1 is $\lambda x[P_1(x)\ \&\ F_1(x)]$, and $M_{Q,C}$ abbreviates the following expression

$$M_{Q,C} = \exists x[C(x)\ \&\ \forall y\ [C(y) \supset (x=y)]\ \&\ Q(x)].$$

Thus far we have considered two types of cross-level references: the trivial one where $L_1 = L_2 = L_0$, and the remote reference in supercontext, that is, with $L_1 < L_2 = L_0$. Let us now turn to the remaining type, for which $L_0 = L_1 < L_2$. We call such situations **remote references in subcontext**.

Example 9. Let TP and T be as in Example 7. This time suppose that we have the following situation: $L_0 = L_1 < L_2 \supseteq L_{+1}^{tp,T}$, where **tp** is TP_t for some $t \in T$. Our discourse may now look as follows.

9a. A president$_1$ sits in the first row.
9b. The president$_2$ is elected every four years.

Sentence 9b has the following translation with a remote reference to "the president" in (9a).

$$9b \rightarrow \exists x\ [p(x)\ \&\ \exists t\ [SL(t)\ \&\ p((x\ t))\ \&\ sfr((x\ t))]$$
$$\&\ \forall y\ [(p(y)\ \&\ \exists t\ [SL(t)\ \&\ sfr((y\ t))]) \supset$$
$$(x{=}y)]\ \&\ eefy(x)]$$

where "sits in the first row" $\rightarrow$ sfr.

The formula translating (9b) says that the president, a unique superobject denoted by x, is elected every four years. The uniqueness of the president superobject is determined by the fact that one of its instances with respect to some space-time coordinate, identified in the formula by the term (x t) where t is an element of this coordinate, has the property of sitting in the first row.

One may wonder whether the use of the definite description in (9b) is always necessary to maintain the remote reference, since we rejected such a necessity in supercontexts. It should be clear enough, considering the level structure introduced earlier, that we can and have to use definite descriptions in subcontextual references. We illustrate this point with another example. The problem is discussed at a great length in Strzalkowski (1986).

Example 10. In the following discourse, the descriptions a president$_1$ and a president$_3$ refer to objects which are not relatively singular and therefore do not belong to the same level. Although the discourse seems connected and coherent, there is no straightforward correspondence between the two sentences.

10a. John wants to become a king rather than a president$_1$.
10b. That is because a president$_3$ is elected every several years, while the king rules for a lifetime.

Here a president$_1$ refers to an L_0-instance P_1 of an L_{+1}-level object P_2 specifically $P_2 \in L_{+1}^{P_1,T}$. A president$_3$ refers to some object P_3, which is still non-singular from the perspective of level L_0. P_3 is an instance of P_2 but with respect to a coordinate T' different than T. In other words, P_3 is placed on the level $L_{-1}^{P_2,T}$, which is not a part of L_0. There is no direct correspondence between a president$_1$ and a president$_3$ beyond the fact that $P_2 \in L_{+1}^{P_1,T} \cap L_{+1}^{P_3,T'}$. It may well happen, however, that for some coordinate U and some $u \in U$, we shall have $(P_3$ $u)$ denoting P_1. This is why John would rather be a king . . . . What relates these two "presidents" is a composition of two remote references: the one made by a president$_1$ to the object P_2 (we may wish to call this object the president$_2$), and the other made by a president$_3$ to P_2. This same effect will be produced if we use other measurably singular descriptions like "every president," "some presidents," or "most presidents" in (10b). If we had the definite "the president" instead of indefinite "a president" in (10b), however, we would obtain a clear instance of a remote reference in subcontext, that is, with $P_2 = P_3$. The objects and their respective levels in the presidents case are illustrated in Figure 3.
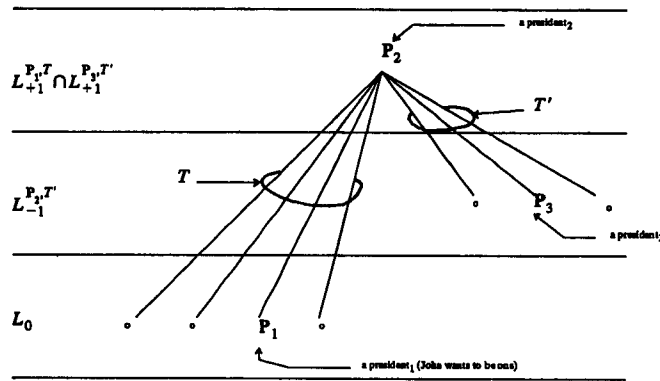
Figure 3. The structure of levels in the president example (Example 10).

The case of the remote reference in subcontext is summarized with the translation Rule 12.

Rule 12 (Subcontext Translation Rule):

If the context-setting sentence S_1 with the translation $\exists x\ [P_1(x)\ \&\ F_1(x)]$ is interpreted at some level $L_{-1}^{\xi,T}$, where ξ is an object satisfying the current sentence S_2 when interpreted at L_0, and S_2 contains an unresolved remote reference P_2, that is,

$$S_2 = \exists x\ [P_2(x)\ \&\ C(x)\ \&\ \forall y\ [(P_2(y)\ \&\ C(y))$$
$$\supset (x=y)]\ \&\ F_2(x)]$$

then the full translation of S_2 is obtained as

$$S_2(P_2) \rightarrow \lambda C[L_2](\lambda u[\exists t\ [C_1((u\ t))]])$$

where the subcontext $C_1 = \lambda x\ [P_1(x)\ \&\ F_1(x)]$ is derived from S_1.

7 Superobjects vs. Plural Terms

We now examine the nature of superobjects, that is, the objects placed at level L_{+1}. In particular, we are interested in what sets them apart from their instances. Let us again consider the term "The President" as referring to the superobject **TP** at some level L_{+1} (cf. Example 7). We may say that such **superterm** has the property of **collective referring** to all of its instances at once, but without necessarily making a reference to any of these instances in particular. At level L_0, if we want to refer to a set of objects, we use one of the enumerating quantifiers "every" or "each," or we use some sort of plural, such as in "These presidents were married" (*each*). We note, however, that most plurals, especially the so-called "bare" plurals, like presidents, tigers, or meetings can actually refer collectively, which makes them akin of superterms, though many of these plurals can be ambiguous between the collective and a non-collective interpretations. Just like in the case of singular superobjects, the objects denoted by plural terms, such as presidents, tigers, Americans, etc., cannot be always identified with the corresponding sets of lower-level instances. It turns out that the plural terms actually

denote superobjects,[11] and therefore they should be interpreted at the same level as respective singular superterms. We will see that the generalization leads naturally to plural terms that may or may not induce equivalent singular superterms. Conversely, a plural equivalent to a singular term may suggest the most natural coordinate to decompose the superobject in its denotation into lower-level instances. When a singular term lacks a plural equivalent, however, we may admit that the object in its denotation is not naturally decomposable and that we are now looking at the bottom-most level in some decomposition hierarchy. A further decomposition may be still possible, but it can only produce objects that will never assume an independent status and will remain recognized only as instances of some more general superobject scattered over that or another coordinate. This phenomenon is characteristic of so-called mass objects and their corresponding mass terms. As an example, consider such nouns as water, gold, or heat. They lack plural equivalents, and there is no obvious way to decompose them into lower-level instances, except with 4–dimensional space-time coordinate. Note that although we can occasionally use a morphologically plural term, like waters, these usually will not be equivalent to singular superterms, and they will often denote other mass objects as well, such as in "waters of the Nile." Space-time coordinates can also be used to obtain instances of otherwise non-decomposable entities, such as John, into space-time slices, called "stages" (Carlson 1982). Quite naturally, the question of where one level ends and another begins arises. The following two examples provide some insight into the level-boundary problem.

Example 11. Consider the following sentences.

11a. Mary brings (some) water every day.

11b. John picks up the mail every morning.

Let "water" in (11a) be the name of some superobject **W** at level L_{+1}. Presumably, Mary brings only a part of **W**, but we can say that **W** is being brought by Mary every day. This is the same **W** every day, although each time possibly a different part of it is in transit, which leads to the obvious translation (at L_{+1}),

i. 11a $\rightarrow$ brings-every-day(M,W)

where M and W are individual constants denoting L_{+1} objects **M** (Mary) and **W** (water). Alternatively, if we evaluate (11a) at level $L_0 = L_{-1}^{W,T}$ with a space-time coordinate T, in which case (11a) should read "Mary brings some water every day," we obtain the following interpretation:

ii. 11a $\rightarrow$ $\forall x\ [day(x) \supset \exists t\ [brings(x, M, (W\ t))]]$

where **W** is as before, **M** is an individual constant denoting a space-time stage of Mary,[12] and brings(x,y,z) should be read *at x y brings z*. Now (W t) denotes some instance of the superobject **W** at level L_0, of which we

can say that it is water as well; that is, water$((W\ t))$. At L_0, we replace $(W\ t)$ by a singular variable u, obtaining

ii'. 11a $\rightarrow \forall x\ [d(x) \supset \exists u\ [water(u)\ \&\ brings(x, M, u)]]$

The translations featured in (i), (ii), and (ii') may appear to be equivalent statements made from different perspectives and at different levels of detail. This is not the case, though. In particular, (ii) and (ii') are not equivalent. This point is further illustrated with (11b). Let **m** be the mail that John is picking up at any one particular occasion. There is a space-time coordinate T such that the level $L_{+1}^{m;T}$ contains a unique superobject for which the following holds.

iii. 11b $\rightarrow \exists x\ [mail(x)\ \&\ C(x)\ \&\ \forall y\ [(mail(y)\ \&\ C(y))$
$\supset (x = y)]\ \&\ picks\text{-}ev\text{-}morn(J, x)]$

The context C is used to determine the uniqueness of John's mail superobject. Note that the quantification ranges over objects at $L_{+1}^{m;T}$. Because John's mail superobject is decomposable with T we obtain the following equivalent cross-level interpretation.

iv. 11b $\rightarrow \exists x\ [mail(x)\ \&\ C(x)\ \&\ \forall y\ [(mail(y)\ \&\ C(y))$
$\supset (x = y)]\ \&\ \forall u\ [morn(u) \supset \exists t\ [picks(u, J, (x\ t))]]]]$

Here, t ranges over elements of coordinate T, and $(x\ t)$ refers to an L_0 instance of John's mail superobject. The three arguments of picks(x,y,z) are understood as *at x y picks up z*. As it stands, (iv) still needs to be interpreted at $L_{+1}^{m;T}$. To reduce this translation to the form that would be interpretable at L_0, we have to replace the definite reference in the first line of (iv), that is,

$$\exists x\ [mail(x)\ \&\ C(x)\ \&\ \forall y[(mail(y)\ \&\ C(y))$$
$$\supset (x = y)]]$$

with a predicate, interpretable at L_0, that would uniquely indicate John's mail. Then replacing $(x\ t)$ by a variable z, we obtain the following formula.

iv'. $\forall u\ [morn(u) \supset \exists z\ [John\text{'}s\text{-}mail(z)\ \&$
$picks(u, J, z)]]$

Note that equivalence to (iv) requires the John's-mail predicate to be true of an empty z, which can happen if, at some occasion, John receives no mail at all. Otherwise (iv) and (iv') are not equivalent. In (iv) we can maintain truth of picks $(u,J,(x\ t))$ because we do not require $(x\ t)$ to denote anything at L_0; we merely say that a t exists. This example actually shows the power of the multilevel representation. Observe that we have just delivered a very strong argument supporting the claim that superobjects are not merely sets of their instances. It must be noted here again (cf. footnote 7) that some limit needs to be set up for the amount of exceptions to a general statement like (iv) that we are willing to tolerate.

There remains one more reading of (11b) that does not seem to require any reference to higher-level objects. This reading could be paraphrased as: Every morning there exists *the* mail such that John picks it up, or more formally

v. 11b $\rightarrow \forall x[morn(x) \supset \exists u\ [mail(u)\ \&\ C(u)\ \&$
$\forall y[(mail(y)\ \&\ C(y)) \supset (u = y)]\ \&\ picks(x,J,u)]$

This time the variable u is bound at L_0. Is this translation feasible? We can answer both "yes" and "no." A "yes" answer indicates that the transformation gives us a singular interpretation of (11b) at L_0. Note that because of the uniqueness clause in it, (v) says no more than that John keeps picking up the same thing every morning, since the context C does not depend on x and is the same each time. A "no" answer indicates that the latter interpretation most probably does not express our intention. The translation of (v), although possible, is not equivalent to either (iii), (iv), or (iv').

At the beginning of this section we stated that bare plurals, like presidents or meetings, behave much like singular superobjects. Let us look somewhat closer at this issue now. Consider the pairs of sentences (12a,b) and (12c,d) below.

Example 12.
12a. The faculty meeting is held every month.
12b. Faculty meetings are held every month.
12c. The rat can live in most countries.
12d. Rats can live in most countries.

It seems to us that the preferred interpretation of sentences in (12b) and (12d) is such that the plural noun phrase (faculty meetings, rats) is understood as denoting a higher-level concept. For the same reason as (12a), and (11b) before it, (12b) will remain a truthful L_{+1} statement in spite of the fact that, at some occasions, meetings can be canceled (due to holidays, for example). In the case of (12c) and (12d) the point is somewhat finer. This time we may be less reluctant to say that (12d) means, in fact, that *every* rat can live in most countries, because we talk about possibilities. There is a catch in here, however; though every rat can live at many different places, not every one, or perhaps even none, can live at most of these places.

One of this section's key observations is that plural terms are, in many respects, equivalent to singular superterms. We can assume, for now, that plural terms, cautiously named prototypes of superterms, actually denote superobjects as well. For some plural terms, we can find equivalent singular versions like tigers—tiger, presidents—president, etc. Others do not have this property. Alternatively, mass terms will usually lack plural equivalents, for example, "water" cannot be identified with "waters" in general. Still others may expose a surprising mixture of the properties, like "people," which is a plural term with morphologically singular form and which may also be used as mass superterm.

8 CONCLUSION

We presented an approach to representing various kinds of non-singular concepts in natural language discourse.[13] The major observation of our theory is that

reality, as perceived by an intelligent individual, can be regarded as a partially ordered structure of levels such that each level contains only those objects that are considered relatively singular. Observe that there are virtually no restrictions imposed upon the notion of relative singularity, so that the distribution of objects between levels of the world model may differ among different individuals. Non-singular objects, called superobjects, are placed at a number of higher levels, which are related to the current level with various coordinates. Conversely, a singular object may be decomposed along a coordinate, and new objects, so obtained, will be placed at some lower level. This same coordinate can be used then to obtain instances of other objects at this lower level, so that the relative singularity of objects within each level is maintained. This theory also contributes to a better understanding of discourse internal cohesion by introducing the notion of remote reference in text.

We believe that our approach promotes computational tractability of some more difficult properties of natural language discourse. Among these, the notion of intension as formalized by Montague (1974a-d) has long been considered difficult for a practical realization, and a more tractable alternative has been sought ever since. Our theory of names and descriptions takes us a step in this direction, although an implementation or a formal complexity evaluation have yet to be attempted. Observe that with the concept of superobject and coordinate, we no longer have to identify general terms like "temperature," "president," or "water" with intensions, that is, functions over possible worlds. Superobjects are not some mysterious, extra-world entities, but they acquire a concrete status which makes them as comprehensible as ordinary objects. In a simplified account, our notion of coordinate can be loosely related to the possible world theory's concept of index. The coordinate is, however, far more selective than the index. In a particular discourse situation, we can pick up that aspect of intensionality of some concept that is, at present, relevant to our understanding of the discourse. Our freedom in selecting that or another coordinate is all important. Note also that structures of coordinates may vary considerably. For example, a pure time coordinate may consist of time points as well as consists of time periods, and time periods are of much greater significance in practice (examine "the president" example). Coordinates connote more than indices. They can be used to define a non-singular status of an object which is otherwise purely extensional (the examples of "water" from the last section). In this sense, the concept of non-singularity has a local, and often subjective, character.

ACKNOWLEDGMENTS

REFERENCES

Barwise, J. and Perry, J. 1983 *Situations and Attitudes*. The MIT Press, Cambridge, MA.

Barwise, J. and Perry, J. 1885 "Shifting situations and shaken attitudes." *Linguistics and Philosophy* 8;(1): 105–161.

Bronnenberg, W.; Bunt, H.; Landsbergen, J.; Scha, R.; Schoenmakers, W.; van Utteren, E. 1980 "The Question Answering System PHLIQA1." In L. Bolc (ed.) *Natural Language Question Answering Systems*. Macmillan.

Carlson, G. 1977 "A unified analysis of English bare plural." *Linguistics and Philosophy* 1: 416–456.

Carlson, G. 1982 Generic Terms and Generic Sentences. *Journal of Philosophical Logic* 11: 145–181.

Donnellan, K. 1971 Reference and Definite Descriptions. In D. D. Steinberg, L. A. Jakobovits (eds.) *Semantics*. Cambridge University Press, Cambridge, England; 100–114.

Kripke, S. 1972 Naming and Necessity. In D. Davison, G. Harman (eds.) *Semantics of Natural Language*. Reidel, Dordrecht; 253–355.

Landsbergen, J. and Scha, R. 1979 "Formal languages for semantic representation." In Allen and Petofi (eds.) *Aspects of Automatized Text Processing: papers in text linguistics*. Buske, Hamburg.

Lewis, D. 1976 General Semantics. In B. H. Partee (ed.) *Montague Grammar*. Academic Press; 1–50.

Montague, R. 1974a On the Nature of Certain Philosophical Entities. In Thomason 1974 *Selected Papers of Richard Montague*. Yale University Press, New Haven, CT.

Montague, R. 1974b English as a Formal Language. In Thomason 1974 *Selected Papers of Richard Montague*. Yale University Press, New Haven, CT.

Montague, R. 1974c Universal Grammar. In Thomason 1974 *Selected Papers of Richard Montague*. Yale University Press, New Haven, CT.

Montague, R. 1974d The Proper Treatment of Quantification in Ordinary English. In Thomason 1974 *Selected Papers of Richard Montague*. Yale University Press, New Haven, CT.

Partee, B. H. 1972 Opacity, Coreference, and Pronouns. In D. Davison, G. Harman (eds.) *Semantics of Natural Language*. Reidel, Dordrecht; 415–441.

Quine, W. V. 1960 *Word and Object*. The MIT Press, Cambridge, MA.

Quine, W. V. 1973 *The Roots of Reference*. Open Court. La Salle, IL.

Strzalkowski, T. 1986. A Theory of Stratified Meaning Representation. Doctoral dissertation, School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada.

Strzalkowski, T. 1986a An Approach to Non-Singular Terms in Discourse. In *Proceedings of the 11th International Conference on Computational Linguistics (COLING-86)*. Bonn, West Germany.

Strzalkowski, T. 1989 "A meaning representation for generic sentences." (unpublished manuscript).

Strzalkowski, T. and Cercone, N. 1985 A Framework for Computing Extra-Sentential References. In *Proceedings of the Theoretical Approaches to Natural Language Understanding*. Halifax, Nova Scotia; 107–116.

Strzalkowski, T. and Cercone, N. 1986 A Framework for Computing Extra-Sentential References. *Computational Intelligence* 2(4): 159–180.

Thomason, R. (ed.) 1974 *Selected Papers of Richard Montague*. Yale University Press, New Haven, CT.

Vendler, Z. 1971 Singular Terms. In D. D. Steinberg, L. A. Jakobovits (eds.) *Semantics*. Cambridge University Press, Cambridge, England; 115–133.

NOTES

1. For Montague's system to work it is necessary that a possible world is understood as a completely specified *alternative reality*, and cannot be thought of as merely a different model, or a partial situation (Montague 1974a, Barwise and Perry 1985). Although various "Montague-style" or "Montague-inspired" approaches to natural language processing have been proposed in computational linguistics (see, for example, Landsbergen and Scha 1979, Bronnenberg et al. 1980) none of them ever attempted to tackle a full-scale semantics for IL.

2. "Profiles: A gentle reign," by Susan Orlean. *The New Yorker*, Dec. 12, 1988, p. 50.

3. *is* in "is rising" and in "is ninety" are of course different uses of the verb *to be*. This is the identity *is* in the second sentence that we are concerned with here.

4. This example was suggested by one of the reviewers.

5. We use **boldface** to denote objects, and romanface to indicate their names and other descriptions referring to them.

6. There is no connection between this and other drawings presented in this paper, and any semantic network formalism.

7. This is not to say that we can have just any amount of negative evidence in the form of literals $\neg fly((B\ t))$ and still keep on believing that fly(B). An empirical verification or refutation of the latter will depend upon the number of instances seen and the quantitative relationship between the positive and the negative evidence (Strzalkowski 1989). Nonetheless, we have to be prepared to accommodate exceptions, up to a point. The discussion of these issues is beyond the scope of the present article.

8. A detailed discussion on how to produce such translations can be found in Strzalkowski and Cercone (1986) and Strzalkowski (1986).

9. This is, quite clearly, due to the temporal aspect of the majority of sentences produced in natural language discourse (use of verb tenses and temporal adverbs).

10. Before a translation rule like one of these presented in this paper can be used, a sentence must undergo numerous transformations within the Stratified Model (Strzalkowski 1986). These transformations include, but are not limited to, morphological analysis, lexical disambiguation, and syntactic parsing.

11. Quine (1960) seems to suggest this conclusion; refer, for example to his double interpretation of mass terms (pp. 120–121).

12. To be precise, we should represent Mary-stage as $(M\ t)$, that is, as a term referring to an instance of the object M. However, our naming convention discussed earlier allows for replacing the definite description $(M\ t)$ by the name of the L_0 object in its denotation. We utilize this option here.

13. This approach to representing non-singular concepts contributes to the transformation F_{n-1} of the Stratified Model; see Strzalkowski (1986). In particular, rules 11 and 12 add greatly to the explanation of text cohesion, about which transformation. F_{n-1} is concerned, although they introduce some ambiguity into the transformation.