

LABORATORY FOR  
COMPUTER SCIENCE



MASSACHUSETTS  
INSTITUTE OF  
TECHNOLOGY

MIT/LCS/TM-101

DESCRIPTIONS AND THE SPECIALIZATION OF CONCEPTS

WILLIAM A. MARTIN

MARCH 1978

MIT/LCS/TM-101

**Descriptions and the Specialization of Concepts**

by

**William A. Martin**

**Laboratory for Computer Science  
Massachusetts Institute of Technology**

### Abstract

The OWL II system computes with expressions which describe an object from a particular viewpoint. These partial descriptions form a tree structure under the specialization operation, which preserves intensional properties. The descriptions are also related in terms of their extensions by characterization and exemplar links. Descriptions of individuals must always specify a context of the description. Eight ways in which one description can be a specialization of another are distinguished.

### Acknowledgement

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Office of Naval Research under contract no. N00014-75-C-0661. I especially wish to acknowledge the informed suggestions and steadfast support of William Carlson, the ARPA contract monitor. I thank Lowell Hawkinson and Peter Szolovits for valuable advice and discussion during the formulation of these ideas. Mitch Marcus, Gretchen Brown, and Kurt VanLehn provided helpful comments on a draft of this paper.

---

Key words: Artificial Intelligence Knowledge Representation Semantics

## 1. Introduction

OWL II is a language for knowledge representation under development at the MIT Laboratory for Computer Science [2]. OWL II is based on the Linguistic Memory System (LMS) developed by Hawkinson [1].

LMS provides two basic data operations, specialization and attachment - these correspond roughly to CONS and PUTPROP in LISP. In this paper problems and phenomena associated with the development of representation (IS-A, AKO, etc) hierarchies are examined. I show how these are dealt with in OWL II, using constructs built on specialization and attachment.

## 2. The Distinction Between a Description and Its Referent

The distinction between a description and its referent has been known to philosophers for many years. As it turns out, it is an important distinction to make in computational linguistics. For example, sentence 1

1. The miner and sapper went to work.
2. The miner and the sapper went to work.

can refer to either one individual or two, depending on whether the conjunction and is taken to conjoin the descriptions miner and sapper or their referents. This ambiguity even exists in 2. Similarly, the phrase old friend can refer either to someone who is old, or someone with whom an old friendship exists, depending on whether old is taken to apply to the referent of friend or the description, friend, itself. (3) So-called "hedges" (23) appear to be meta-level instructions on how a description should be used in finding a referent. For example, in fake x, fake says that the referent is something which has properties allowing it to be recognized as an x, but which fails to meet some of the required criteria of an x. Note particularly that fake x's are not x's which are fake.



Sentence 3

3. Esther Williams is a regular fish.

shows that regular specifies that the referent has some of the functions but not the form of the following description.

These examples show clearly the distinction between computing with a description or its referent.

### 3. Specialization

Given a description, it makes sense to talk about the set of individuals to which that description applies. These individuals will be termed potential-referents. For example, while a dog has a single referent, it has many potential-referents. A description may or may not have a unique potential-referent. It may be inherently non-referential, only applying to other descriptions, e.g., very; it may have no potential-referents, e.g., a round square; it may have one, e.g., the President of the United States in 1978; or it may have many, e.g., a dog. Given several descriptions one can form the set of individuals to which they all apply simultaneously. Given some finite set of descriptions and some finite set of individuals, one could in principle form the potential-referent sets  $R_i$  which satisfied each subset  $S_i$  of the descriptions. For example, given the descriptions dog, cat, barker, and housepet, and the individuals in the real world, the non-empty sets  $R_i$  would correspond to the subsets of descriptions {dog} {cat} {barker} {housepet} {dog barker} {dog housepet} {cat housepet} {barker housepet} {dog barker housepet}.

In OWL II, descriptions are arranged in a tree structure as shown in Figure 1. This tree structure serves as a basis both for storing and retrieving descriptions and for inheritance. The

use of simple tree structures has been rejected in other knowledge representation schemes because it is obvious that one individual can be, for example, simultaneously a puppy, a biter, and a pet. But this is not necessarily a constraint on the hierarchy of the descriptions, rather it is a constraint on the hierarchy of the potential-referent sets,  $R_i$ , of individuals satisfying descriptions. In OWL II, each description is represented explicitly as a node in a semantic net. A potential-referent set,  $R_i$ , is represented only indirectly as links between its corresponding descriptions.

In OWL II, we let semantic memory be made up of concepts and symbols. Symbols are written as character strings between double quote marks, e.g., "ENGINE". As shown in Figure 1, the most general concept is SUMMUM-GENUS. Symbols are taken to be atomic in the sense that they cannot be decomposed in any way. Concepts are non-atomic. They are constructed from SUMMUM-GENUS and symbols by using the binary operation, specialization.

Specialization is written:

(genus specializer)

where genus  
is a concept and  
specializer is a  
concept or symbol.

We say that a concept is a specialization of the concept in its genus position.

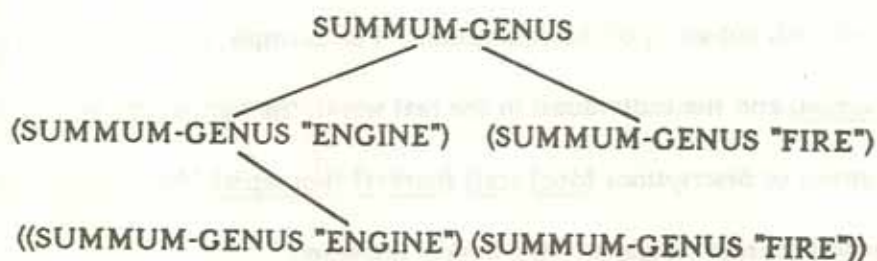


Figure 1

For example in Figure 1 we have constructed two specializations of SUMMUM-GENUS:

(SUMMUM-GENUS "ENGINE") and (SUMMUM-GENUS "FIRE"). We have then specialized (SUMMUM-GENUS "ENGINE") by (SUMMUM-GENUS "FIRE").

Moving up in the genus direction, it is clear that concepts are the nodes of a tree with SUMMUM-GENUS at the root. SUMMUM-GENUS is taken as a specialization of itself;

SUMMUM-GENUS = (SUMMUM-GENUS "SUMMUM-GENUS")

We say that any concept, C, forms a class which contains all the concepts in the sub-tree whose root is C, including C itself.

If specialization is carried to very many levels, the expression for a concept quickly becomes unwieldy. We avoid this through the familiar mechanism of labeling. The expression

label = concept

where label is any string of letters  
digits, hyphens, and periods.

assigns label to concept. A label is just a notational abbreviation for the parenthesized expression that exhibits the genus and specializer of a concept; it has no semantic significance in and of itself.

Using labels we might rewrite Figure 1 as Figure 2.




Figure 2

The phenomenon we model with specialization is called syntagma by Marchand (4) in his study of compounds. Explaining it, he states

"2.1.1 The coining of new words proceeds by way of combining linguistic elements on the basis of a determinant/determinatum relationship called syntagma. When two or



more words are combined into a morphological unit on the basis just stated, we speak of a compound. In the system of languages to which English belongs the determinant generally precedes the determinatum. The types which do not conform to this principle are either syntactical compounds (e.g. *father-in-law*) or loan compounds (e.g. *MacDonald, Fitzgerald*) with the 'inner form' of a non-English language.

2.1.2 The principle of combining two words arises from the natural human tendency to see a thing identical with another one already existing and at the same time different from it. If we take the word *steamboat*, for instance, identity is expressed by the basis *boat*, the difference by the word *steam*. *Steamboat* as compared with *boat* is a modified, expanded version of *boat* with its range of usage restricted (see below) so that *steamboat*, the syntagma, will be found in basically the same semantic contexts as the unexpanded *boat*. The syntagma *steamboat* also retains the syntactic primary feature of *boat*, *steamboat* belongs to the same word class 'substantive' to which *boat* belongs. An adjective such as *color-blind* is an expansion of *blind*. A person is called *color-blind* because he is basically seen as *blind* though only so with regard to colors. *Rewrite* as compared with *write* is basically the verb *write* with which it is to a great extent exchangeable except for the modification expressed by *re-*. This does not, however, affect the word class of the syntagma, which is that of a verb.

Combinations of types *steamboat*, *colorblind*, and *rewrite* which are mere morphological extensions of the words *boat*, *blind*, and *write* respectively, will be termed EXPANSIONS. An expansion will then be defined as a combination AB in which B is a free morpheme (word) and which is analysable on the basis of the formula  $AB = B$ . This means that AB belongs to the same word class and lexical class to which B belongs. Combinations of the kind illustrated by *steamboat* and *colorblind* which contain free morphemes both for the determinant and the determinatum will be termed compounds. Combinations of the type *rewrite* where the determinatum is a free morpheme while the determinant is a bound morpheme are prefixed words. Both compounds and prefixed words thus are subgroups of the larger class called 'expansions'.

2.1.3.1 A further clarification may not be out of place. Semantically speaking, the determinatum represents the element whose range of applicability is limited by the determinant. A *steamboat* is basically a *boat*. But whereas *boat* as an independent unit can be used with reference to an unlimited variety of boats, the applicability of *steamboat* is limited to those which are powered by steam, excluding those which are not steamboats. We might say that this exclusion in *steamboat* of 'non-steamboat' things constitutes the determination of *boat* as performed by the first element *steam*, which has therefore been called the determinant. *Boat*, as the element undergoing a semantic restriction or determination, has been called the determinatum. However, as a syntagma is a grammatical, not a semantic entity, we would say that the terms determinatum and determinant should be defined as grammatical terms. Grammatically speaking, the determinatum is that element of the syntagma which is



dominant in that it can stand for the whole syntagma in all positions, as has just been stated in a formula.

2.1.3.2. It is important to stress the grammatical character of a syntagma. Semantically speaking, the grammatical determinant is in many cases the part that can stand for the whole combination. This would first apply to compounds of the type *girl friend*. *Girl* may well fill the place of *girl friend*, but it has not become the grammatically dominant part. The semantic dominance of the determinant over the determinatum is, however, most in evidence in derivation containing an appreciative suffix, as in *streamlet* 'little stream'. A *streamlet* is basically a *stream* though an (emotionally) small one, and could therefore take the place of *stream*, if semantic considerations were the criterion of substitution. A *blackish suit* could substitute for a *black suit* as from a purely semantic point view *black* has merely been expanded into *blackish*. But grammatically speaking, *black* in *blackish* has lost its independence to *-ish* just as in *blacken* it has lost its independence to *-en*. In either case it is the suffix that dominates grammatically."

In sections to follow the notion of specialization will be further refined. But first, attachment will be introduced, lest the reader begin to feel that everything must be solved with specialization alone.

#### 4. Attachment

In OWL II each concept has a reference area. Concepts are placed in the reference area of a given concept using attachment. Attachment can be denoted using a complex like that in Figure 3.

```
[ PROFESSOR      #CHARACTERIZATION FACULTY-MEMBER
                  #EXEMPLAR ASSISTANT-PROFESSOR ASSOCIATE-PROFESSOR
                  #PREDICATE ABSENT-MINDED ]
```

Figure 3

Example of a Complex

Symbols prefixed with "\*" are the names of zone relations. Zone relations differ from the arc labels of most semantic nets only in that they belong to the meta-language of OWL II rather than to the knowledge being represented. The leftmost element of a complex is termed the subject. This is

followed by a zone relation, the concepts in that zone, another zone relation, the concepts in that zone, etc.

We now turn our attention to the three most important zone relations:

•CHARACTERIZATION, •EXEMPLAR, and •PREDICATE.

## 5. Characterizations and Exemplars

We may say of dogs that they are quadrupeds, pets, and barkers. The view taken in OWL II, and some other knowledge representation languages (5), is that the knowledge of any individual or class of individuals consists entirely of a set of such descriptions in terms of which questions about the individuals may be answered. The hoped for advantage of this approach is computational efficiency and elegance. Upon learning, for example, that an individual is an elephant, it is not necessary to copy all of the knowledge of elephants onto the individual before answering questions about his color, size, etc. Rather, the view is taken that until proven otherwise "if you've seen one elephant you've seen them all". Questions about the individual are referred to the description of elephants.

Questions of how to proceed when an individual is characterized by a set of such descriptions are, as yet, largely unanswered. First, a decision must be made as to whether the machine will proceed as if the descriptions for an individual are merged into one conglomerate description or whether they will be kept distinct. Keeping descriptions distinct means there must be some explicit means for the programmer to control which description is investigated when trying to answer a particular question. The decision to merge has been made, for example, in FRL (6) and it is the standard mode of operation in NETL (7). Descriptions are kept distinct in OWL II and KRL.

Each description can be considered a representation of the individual from a different point of view. One needs to know from what points of view a particular question might be answered. Smith (8) has suggested that descriptions could be categorized into basic categories such as form, function, purpose, etc. However, this would rule out the use of descriptions such as "nuclear powered attack submarine" which constrain the form, function, and purpose simultaneously. A different approach can be seen in work by Long (9) and Sussman (10). These authors define rather global viewpoints from which a problem such as program writing or circuit analysis can be attacked and then categorize descriptions according to global viewpoint.

Suppose that a DOG is a QUADRUPED, PET, and BARKER. A Venn Diagram of the intersection of the potential-referent sets of these descriptions is shown in Figure 4. In general, things will be known about a DOG which are further refinements of what is known about a QUADRUPED. For example every QUADRUPED employs a series of gaits, but a DOG employs a particular series. Every QUADRUPED has feet, but the feet of a DOG may be further structurally described to have toes. Similar refinements can be made for a dog as a PET or a BARKER. Let DOG-QUADRUPED, DOG-PET, and DOG-BARKER be the labels for these refined descriptions.




Figure 4




We have to determine the relationship between DOG, DOG-QUADRUPED, DOG-PET, DOG-BARKER, QUADRUPED, PET, and BARKER. In OWL II we take DOG-QUADRUPED, DOG-PET, and DOG-BARKER, to be specializations of QUADRUPED, PET, and BARKER, respectively. Specializations inherit description from their genus unless explicitly overridden. Next, common parts of the descriptions are identified and linked. For example, every BARKER must have a MOUTH, but the MOUTH of a DOG-BARKER is the MOUTH of DOG-QUADRUPED.

We come now to the relationship between DOG, and the combination DOG-QUADRUPED, DOG-PET, DOG-BARKER. In KRL-0, DOG is this combination. In OWL II, DOG must be chosen to label one of these; the obvious choice being DOG-QUADRUPED. Note that the same combination DOG-QUADRUPED, DOG-PET, DOG-BARKER is formed in either case, but while in KRL this combination is an explicit structure of the language, it exists only implicitly in OWL II. The choice affects:

- a) Where in this combination DOG will point and how its parts will be addressed from there.
- b) How one moves from one of these viewpoints to another.
- c) Where new facts about DOG are placed and how they relate to the existing description.

The potential-referent sets of two descriptions may be mutually exclusive, may intersect, or one may contain the other. For example, dog and cat are mutually exclusive, dog and pet intersect, while all dogs are barkers - with the exception of barkless besenji's and cripples. These relationships are indicated in OWL II by defining a \*CHARACTERIZATION of a concept, C, to

be a concept whose potential-referent set includes or is equal to the potential-referent set of C. An **EXEMPLAR** of a concept C is a concept whose potential-referent set is included in the potential-referent set of C. This is illustrated by Figure 5.



A is a #CHARACTERIZATION of B

C is an #EXEMPLAR of B

Figure 5

Venn Diagram showing the relationship between the referent set of a concept and those of its characterizations and exemplars.

In the case of intersecting sets like those of DOG and PET, the concept DOG-PET is defined to describe the intersection. The above dog example becomes

[DOG #CHARACTERIZATION DOG-BARKER #EXEMPLAR DOG-PET]

Recall that specializations inherit the description of their genus, unless that description is explicitly overridden. This allows besenji's to be described as dogs which don't bark.

It is interesting to note that if A is a characterization of B then B implies A. Consequently, programs can be written to access characterizations and exemplars to do antecedent and consequent reasoning as in Micro-Planner (14).

## 6. Slots and Predication

It is convenient to think of OWL II concepts as having slots (11). For example, the concept HIT might have, among others, the slots (SUBJECT HIT) and (OBJECT HIT), representing the two main roles in a hitting. When an instance of hitting takes place, the roles are instantiated with particular actors and these actors play out some version of the scenario of hitting.

We define a predicate to be a concept with a SUBJECT slot. In addition to describing a concept with CHARACTERIZATION's and EXEMPLAR's we can also apply PREDICATE's to it. For example, if CHASE-STICK is the predicate CHASE with its OBJECT slot filled with STICK, and BARK is a predicate with only a SUBJECT slot, we could form

[DOG #PREDICATE BARK CHASE-STICK]

which indicates that DOG is the SUBJECT of BARK and CHASE-STICK. The distinction between characterizing a dog as a barker and predicating him with bark has no parallel in logic, where both are treated as predicates. The distinction has been made in OWL II for several reasons.

a) It allows us to distinguish between further description of a viewpoint (PREDICATE) and a change of viewpoint (CHARACTERIZATION). To say that a dog is a barker is to indicate that his barking serves to identify him as appropriate to play other roles. For examples, barkers don't make good neighbors but make good caretakers. We can further describe how to deal with barkers, etc. Just to say that a dog barks is not to indicate that a whole viewpoint of him should be built around this action.

b) Using both #PREDICATE and #CHARACTERIZATION gives us a way to represent distinctions seen in English. We have distinct representations for

Fido barks.	#PREDICATE
Fido is a barker.	#CHARACTERIZATION

Moreover, we can distinguish the relationships between subject and complement in

John is to love her.	#PREDICATE
To see her is to love her.	#CHARACTERIZATION

c) There are implementation issues. For examples, predicates such as red, heavy, etc. can be treated as features, while characterizations form a type heirarchy. It is useful to have both of these mechanisms in matching.

In OWL II ambiguities like old friend are resolved by classifying all predicates into those



which apply to the referents of concepts and those which apply to the concepts themselves. A predicate can be tested to determine its class. In old friend, old has two senses, one a predicate which applies to the referent, one a predicate which applies to the concept, friend.

We have now described a number of constructs for a knowledge representation language. One construct common to such languages which we have not yet discussed is a formal notion of context. Context is best approached by first considering the problems involved in the identification of individuals.

## 7. The Identification of Individuals

As Strawson (12) has pointed out, we think of the world as containing particular things, some of which are independent of ourselves; we think of the world's history as made up of particular episodes in which we may or may not have a part. While from a logical point of view one may argue that a person's feelings and sensory input are more real than objective particulars, it is clear that people nevertheless organize their thoughts with objective particulars as the primary embodiment of reality. Furthermore, a person's application of his knowledge is based heavily on the ability to reidentify a particular as the same individual that he saw before.

In a system such as OWL II which is based on descriptions, one needs to know that a description is adequate to identify an individual. In practice, more identification is required for getting a passport than for cashing a check. One is willing to balance the costs of misidentification with the costs of obtaining more evidence. It seems plausible, though, that this is a special sort of behavior invoked only when the costs of misidentification are known to be high - when one is buying an old painting or something. If a colleague substitutes his copy of a new textbook for

yours, misidentification is almost certain. We are led, then, to mark descriptions which are normally adequate to identify an individual. This was one of the features in a system developed by Brown (24).

In OWL II individuals are represented solely in terms of one or more identifying descriptions. If two descriptions represent the same individual, then each must be a characterization of the other. In the traditional terms of logic, these descriptions are equal because they have the same extension.

One may ask of two descriptions whether they differ in such a way that they can not represent the same individual. For example, "the red block on the table today" and "the red block in the box yesterday" can represent the same individual because we believe that a block's identity as an individual is not changed by a change in location. By contrast, if I make some cloth into a dress, then one individual, the cloth, has changed into another, the dress. While the distinction is clear at the extremes, it is not possible to specify what accumulation of change should constitute change to a separate individual. For example, a horse may have his parts replaced systematically with those of a cow without one being able to say precisely when he becomes a cow.

Most descriptions of individuals do have elements which cannot be changed. For example, one cannot change the number of sides in the description of a square and have it remain a description of the same individual. An event of Bob sneezing cannot be described as someone else sneezing. It is essential that Bob do it for it to remain the same event. Whenever an essential element of description is changed, the same individual is no longer represented. In OWL II we will mark certain elements to be essential; while realizing that a combination of unmarked elements can also represent an essential change.



In the typical case, when a description is further specified the resulting description applies to a subset of the individuals specified by the original, e.g., the set of big dogs is a subset of the set of dogs. But in the case of a "hedge" like fake, the resulting set is mutually exclusive of the original. A fake gun is not a gun, by definition of fake. The set of guns and the set of fake guns have no members in common. A fake x is something which answers to the description of an x, except in some essential part. A question arises as to whether the description fake x should be considered a specialization of the description x. We have said that a specialization inherits the description of its genus, unless this is explicitly overridden. In the case of fake, we know that the description is to be inherited except that it is to be overridden in some unspecified but essential way so that the two sets of referents are mutually exclusive. If we take fake x as a specialization of x, it will inherit the description x, as desired, but we will have the situation where in general no relationship will exist between the sets of referents of a concept and its genus.

Direction in this question may again be taken from Strawson (13). He distinguishes between

- a) a description
- b) the use of a description

and consequently between

- a) the meaning of a description
- b) the meaning of the use of a description

To give the meaning of an expression is "to give general directions for its use to refer to or mention particular objects or persons" e.g. the meaning of this specifies the conditions for its use - the speaker intends the item referred to to remain central to discussion, etc. By contrast, the meaning of a use of this is the particular item referred to in that use. To Strawson, the meaning of a description is intensional, the meaning of the use of a description is extensional. Note that a



description can exist "out of context" but the use of a description must always take place in some context.

Viewed in these terms, we can take the description (GUN FAKE) to be a specialization of GUN and say that (GUN FAKE) inherits intensional meaning from GUN. It makes sense to use the expression fake gun in the same sentences where gun is used - he threatened her with a fake gun, he fired a fake gun.

•CHARACTERIZATION and •EXEMPLAR are defined in terms of the extensions of descriptions, specialization in terms of their intensions and the identification of concepts in memory. In general, a concept and its genus have description in common, a concept and its characterizations have one or more individuals in common.

Strawson (12) distinguishes between relative identification and identification within history.

He says

"A speaker tells a story which he claims to be factual. It begins: 'A man and a boy were standing by a fountain', and it continues: 'The man had a drink.' Shall we say that the hearer knows which or what particular is being referred to by the subject-expression in the second sentence? We might say so. For a certain range of two particulars, the words 'the man' serve to distinguish the one being referred to, by means of a description which applies only to him. But though this is, in a weak sense, a case of identification, I shall call it story relative, or for short, a relative identification. For it is identification only relative to a range of particulars (a range of two members) which is itself identified only as the range of particulars being talked about by the speaker" ... "It is identification within his story; but not identification within history.

Discussing what would constitute identification within history he says "A sufficient, but not necessary, condition ... is that the hearer can pick out by sight or hearing or touch, or can otherwise sensibly discriminate, the particular being referred to, knowing that it is that particular."

Another way to look at this, which will be adopted here, is that any description is only

adequate to identify a particular relative to a given context. Viewing the context as part of the description, we may say that every description of a particular contains an essential element which may be viewed as a context. An individual is identified in a story by a different description from that used to identify the individual in the real world. "Identification within history" means identification of an individual relative to the real world. If the above story continued "the man drinking was Bob Smith", Bob Smith being known to the hearer, then the hearer could mutually characterize his story description and his Bob Smith description. In standard logical terms, he could make these two descriptions equal.

Consider now a circuit-description of some particular variety of radio, and an individual resistor, named R23, that appears in that circuit. (7) The description R23 identifies a particular resistor in the context of that circuit. If we ask which individual resistor in the circuit dissipates the most power, R23 would be a perfectly acceptable answer. R23 is clearly an individual in that circuit. Now each radio we make from this plan will have its own version of resistor R23. These R23's are individuals with respect to the real world radios which contain them.

There is an important distinction between this radio example and the example of the drinking man. In the latter we set the descriptions {the man, story} and {Bob Smith, real world} equal, indicating they identified the same individual. In the radio example it would be wrong to set equal {R23, radio1}, {R23, circuit-diagram}, and {R23, radio2} because the descriptions of R23 in two different radios don't represent the same individual. In this case, the R23's in each radio are specializations of the R23 in the circuit diagram.

From this we see that a concept need not have the same context as its genus. Referring to Strawson's distinction between a description and the use of a description, the use of a description



may be represented by specializing the description by a suitable context to render it unique for referent identification.

It may be helpful to consider another example. Suppose that LISP is written in machine language and that OWL II is written in LISP. Each execution of OWL II is clearly an individual, an individual program execution. That individual can be described as the execution of a machine code program, the execution of a LISP program, or the execution of an OWL II program. It admits of three descriptions, each at a different level of detail. Now suppose (A B) is typed in. This may be described as a string of characters "(", "A", " ", "B", ")". This description identifies the individual typed in in the context of the machine language level description of the individual execution. (A B) may also be described as a LISP list notation. This description identifies it in the context of the LISP description. Similarly, (A B) can be described as an OWL II concept notation and thus identified in the context of the OWL II description.

Choice of conventions for specification of context in a programming system seems a difficult problem. Three of the questions raised are:

- a) Is there any distinction between data structures used as contexts and other data structures?
- b) Is there more than one kind of pointer to context?
- c) Is context specified separately from other descriptive information?

Some systems using context in a manner similar to that proposed are CONNIVER (14), CSAW (15), NETL (7), and Hendrix's partitioned semantic networks (16). Of these, CONNIVER and CSAW use special data structures for contexts. In CONNIVER, data items are partitioned to a given context and simultaneously described as being present or absent in that context. It has not been



demonstrated that treating contexts differently than other data structures has any particular advantages. It has the disadvantage that while contexts often correspond to events, places, etc. they are treated differently.

Partitioned semantic networks are a simple approach to context. Every node has a context pointer to a node representing its context. Every context has pointers to the nodes in it. All the nodes involved in the description of, for example, an elephant, would reside in the elephant context. When mention of an elephant is followed by mention of a trunk, one can then look for a trunk node in the elephant context.

In NETL, Fahlman divided nodes into those representing individuals and those representing types. Only individuals have a context pointer and they must have one. If the context pointer of node A points to an individual node B, then A is thought of as in B. If it points to a type node C, it may be marked as either an in or an of context pointer. If something can be thought of as in or of its context, it is given an in pointer, e.g. the motor in/of a car is given an in pointer, but the mother of a boy is given an of pointer. If a pointer is marked as an of pointer the context does not necessarily determine the location. Otherwise, it does.

Fahlman's scheme agrees with OWL II in requiring that all descriptions of individuals be with respect to a context. His stipulation that context can only be assigned to individuals seems unnecessary. It is not an important issue for Fahlman because he has chosen to associate with each type node an individual node representing the set of individuals of that type. This individual node has a context which is taken as the context of the associated type node.

By classifying context pointers as in or of pointers Fahlman has opted to combine context specification with other descriptive information, but if one is to do this a more refined scheme is in

order. The apparent purpose of the in/of distinction is to indicate whether the context also gives the location. But the paint on the house is not the same as the paint in the house, yet the context does give the location. Saying the paint of a house sounds odd and doesn't mean just that which is on it. This argues for expanding in to include other location prepositions. Similarly, of could be expanded. To call something a reason for an action is to describe its role with respect to that action. Yet the reason is not located in the action and we don't say the reason of an action.

Fahlman's spatial contexts, called areas, are presented as well defined individuals. For example, he would say that elephants exist in Africa. The extent of Africa is well defined and the statement gives no information about how the elephants are distributed in Africa. In many cases the extent of a context is not well defined, one can specify only its center or a prototype (as Fahlman does do with the context of roles). For example, the description malt may be used to describe a certain drink made without ice cream in a region centered at Boston, but without clearly defined boundaries. Outside that area it refers to a different drink made with ice cream.

All of the systems mentioned specify context with a single pointer. This can force the computation of a context whenever a description is to be stored. For example, in Fahlman's system an event would be an individual and would thus be placed in an area. If Fahlman is told only the actors in the event he must choose the area based on the contexts of the actors. Presumably the event, "Bob called Clyde in Africa from Boston" exists in the Occident, while the event Bob started to send a package to Clyde in Africa from Boston exists only in Boston. But, of course, if Bob is at MIT one could say that the event exists only in the area of MIT. Since context is used to control searches, the choice of context and the decision of how much effort to spend on making a context, affect the efficiency of the resulting searches. In the above systems, the decision as to what



constitutes the context of an event must be made when the event is stored, not when it is used.

This means that interpretation of what constitutes the context cannot be made in light of the use.

This is a familiar problem with any uniform indexing scheme used to feed data to a simple search strategy.

The above considerations argue for treating context no differently than other description.

Certain elements of the description would be marked as essential because they establish the context.

A description which did not provide enough information to establish context could not be marked as describing an individual. For example, whether a description describes an individual as being in the real world or in a plan would be taken as an essential element of the description. This is the approach taken in OWL II.

To close this section, some final insight into the need for multiple descriptions of the same individual may be gained by considering three dimensions along which descriptions may differ.

Strawson (13) says:

- (1) They differ in the extent to which the reference they are used to make is dependent on the context of their utterance. Words like 'I' and 'it' stand at one end of this scale -- the end of maximum dependence -- and phrases like 'the author of *Waverly*' and 'the eighteenth king of France' at the other.
- (2) They differ in the degree of 'descriptive meaning' they possess: by 'descriptive meaning' I intend 'conventional limitation, in application, to things of a certain general kind, or possessing certain general characteristics'. At one end of this scale stand the proper names we most commonly use in ordinary discourse; men, dogs, and motor-bicycles may be called 'Horace'. The pure name has no descriptive meaning (except such as it may acquire *as a result of* some one of its uses as a name). A word like 'he' has minimal descriptive meaning, but has some. Substantial phrases like 'the round table' have the maximum descriptive meaning. An interesting intermediate position is occupied by 'impure' proper names like 'The Round Table' -- substantial phrases which have grown capital letters.
- (3) Finally, they may be divided into the following two classes: (i) those of which the



correct referring use is regulated by some *general* referring-cum-ascriptive conventions; (ii) those of which the correct referring use is regulated by no general conventions, either of the contextual or the ascriptive kind, but by conventions which are *ad hoc* for each particular use (though not for each particular utterance). To the first class belong both pronouns (which have the least descriptive meaning) and substantival phrases (which have the most). To the second class belong, roughly speaking, the most familiar kind of proper names. Ignorance of a man's name is not ignorance of the language. This is why we do not speak of the meaning of proper names. (But it won't do to say they are meaningless.) Again an intermediate position is occupied by such phrases as 'The Old Pretender'. Only an old pretender may be so referred to; but to know which old pretender is not to know a general, but an *ad hoc* convention.

A proper name or pronoun may be used to characterize other descriptions of an individual. Almost its sole use is to find these other concepts. It adds little to the stock of knowledge about the individual. In the case of pronouns the concept the speaker intended to characterize is found by applying general rules for the construction of a context for the pronoun. In the case of proper names, the concept must already be characterized by that name in the listener's head.

## 8. Denotation and Opaque Operators

The view of descriptions proposed above relates in an interesting way to some recent ideas of McCarthy (17). By way of introduction, McCarthy points out that from the statements

```
knows(pat, combination(safe1))  
combination(safe1)="45-25-17"  
combination(safe2)="45-25-17"
```

one can derive `knows(pat, combination(safe2))`, which may not be true, by substitution of equal expressions. The standard way of viewing this problem is to say that knows is an opaque operator which blocks substitution in its second argument. Following an approach somewhat like OWL II, McCarthy suggests treating concepts as objects distinct from the objects they denote. Introducing Safe1 as the concept of safe1 and Combination as the concept of combination he writes

```
knows(pat, Combination(Safe1))
```

to assert that Pat knows the combination of safel. The previous trouble is then avoided by taking

$$\text{Combination}(\text{Safe1}) \neq \text{Combination}(\text{Safe2})$$

which McCarthy feels to be reasonable, since we do not consider the concept of the combination of safel to be the same as the concept of the combination of safe2, even if the combinations themselves are the same. The relation between concepts and the objects they denote is given by the denotation function (or partial denotation function)

$$\text{safel} = \text{den}(\text{Safe1})$$

The functions combination and Combination are related in a way which may be called extensional

$$(\forall S)(\text{combination}(\text{den}(S)) = \text{den}(\text{Combination}(S))).$$

This relation allows us to find the denotation of Combination(Safel) in order to open the safe. But since the Knows predicate lacks this extensionality in its second argument, the undesired evaluation is not possible.

McCarthy's formulation is summed up in Figure 6. His treatment rests two plausible assumptions.

Combination(Safel) = concept of the combination of the safe

↓ denotes

45-25-17

Safel names → the concept of the safe

↓ denotes

the safe

safel names →

Figure 6


- a) It makes sense to identify some object as the denotation of a concept. That is, the thing it stands for.
- b) Just because two concepts denote the same thing it doesn't mean they are the same object because they may have different intensional meaning.

Surprisingly, the OWL II formulation of this problem gives the same basic mechanism as McCarthy's, but with a different interpretation.

In OWL II, (COMBINATION SAFE1) and 45-25-17 are both taken to be descriptions of the combination of safe1. Since 45-25-17 also describes the combination of safe2, 45-25-17 is not equal to (COMBINATION SAFE1), but is a CHARACTERIZATION of it. Similarly, 45-25-17 is a CHARACTERIZATION of (COMBINATION SAFE2). 45-25-17 represents an individual only in the context of all strings of two digit numbers. Thus the question of substitution of equals does not arise. A value function may be defined which maps a description D into one of its characterizations, or into a characterization of an exemplar of D, or an exemplar of an exemplar of D, etc. For example,

$$\text{VALUE } \{ (\text{COMBINATION SAFE1}) \} = 45-25-17$$

The OWL II formulation is summarized by Figure 7. It rests on the assumptions:



$$\text{VALUE } \{ (\text{COMBINATION SAFE1}) \} = 45-25-17$$

Figure 7

- a) It makes sense to identify one description as the value of another. Presumably the description used as value is more useful for the purposes at hand - e.g. the digit string can be decomposed and the safe then dialed.



- b) Just because two descriptions describe the same thing it doesn't mean that they are the same object because they may have different intensional meanings.

The reader can compare these assumptions with McCarthy's. The result of a value function would be the characterization of an exemplar, for example, if the description D described a program variable, which was a matrix, and the value described a particular 2x2 matrix by giving its elements. The exemplar of D would be a description of the 2x2 matrix as a program variable and the characterization of this would be the listing of its elements. The definition of the value function may be used as the basis according to which an interpreter finds the value of a description. For example, letting combination be a function which is the value of COMBINATION, and SAFE1 evaluate to itself:

```
value{(COMBINATION SAFE1)} =  
  value{COMBINATION} {value{SAFE1}} =  
    combination{SAFE1} =  
      45-25-17
```

In reasoning about the knowledge of others, like Pat, it may be useful to define a function which replaces descriptions by their characterizations. In doing this we may distinguish between our knowledge, Pat's knowledge, and knowledge we have in common with Pat. If we both know that safe1 is the-blue-safe then we may substitute to conclude that Pat knows the combination of the-blue-safe, assuming he can find all conclusions resulting from knowledge he has. If only we know that safe1 is the-blue-safe, this substitution may still be made, but since the-blue-safe describes the individual only in the context of our knowledge, this expression does not allow Pat to find the combination of the-blue-safe. This expression is not characterized in his knowledge by either SAFE1 or 45-25-17.

The above suggests that reasoning about knowledge can be done merely by placing each

concept in the context of whose knowledge it is. However, Moore (22) points out that this doesn't allow for expressing

Pat knows either that dogs bark or that cats meow.

which is not the same as


Pat knows that either dogs bark or cats meow.

Such statements apparently require a meta-level statement of Pat's knowledge.

### 9. Expressions and Semantic Nets

Many existing knowledge based systems, including the largest and perhaps most successful, MACSYMA, use symbolic expressions, not semantic nets, as their primary data structure. At the same time, many new experimental systems use semantic nets. The LMS system used to implement OWL II provides a data structure which is a semantic net with expressions for nodes. The nodes are created by specialization and the net links are created by attachment. By relying primarily on specialization or attachment respectively, one can create a data structure which is primarily a tree of expressions, or primarily a semantic net. The issue arises as to whether it is better to express information primarily in the links or in the nodes.

In discussing some of the difficulties commonly encountered in semantic nets, Woods (18) points out that if one renders Bob is 7 ft tall in a semantic net as



he is presumably using the subject Bob, and attribute value, 7-feet, as nodes and the attribute,

height, as a link between them. Now Bob is very tall does not give us the attribute value but only a predicate on it; forcing the above scheme to be complicated in some way. The problem is even worse for Bob's height is greater than Mary's which seems most intuitively to be represented by an expression such as

height(Bob) > height(Mary).

One way to look at the problem is that height above is a link name, and the standard semantic net only allows the description of nodes, not links. Another way to see the problem is that in the standard notation there is no node corresponding to the height of Bob, only one corresponding to its value. Thus height of Bob can't be described. Woods suggests the introduction of a node representing Bob's height. That is, in OWL II terms, Bob's height should be represented by an expression rather than a link in a semantic net. Doing this in OWL II we could have

```
[ (HEIGHT BOB) #CHARACTERIZATION 7-FEET
  #PREDICATE VERY-GREAT
  (GREATER-THAN (HEIGHT SUE)) ]
```

Notice that representing Bob's height by (HEIGHT BOB) implies a semantic relationship between the genus, HEIGHT, and the specializer, BOB. In OWL I an attempt was made to specify rules for deducing the semantic relationship of the genus and specializer of a concept from their semantic classes. This proved unsatisfactory. Consequently, OWL II uses meta-specializers to make this relationship explicit. Meta-specializers are described in the following section.

Hendrix (16) attempts to solve Wood's problem by restricting links to be "case slot names", claiming that one rarely wants to describe a case slot. We have found it difficult to distinguish what should be a case slot and what should be a relationship, like father. For example, what is purpose. Thus in OWL II we have restricted our equivalent of link names even more than Hendrix, allowing only the zone relations and meta-specializers.



Hawkinson points out that if BOB is a HUMAN, then it is desirable for (LEG BOB) to inherit properties from (LEG HUMAN). Hawkinson effects this in LMS (the system in which OWL II is implemented) by automatically replacing (A B) with ((A C) B) whenever B is in class C. (A C) is then called the generalizer of ((A C) B). Thus (LEG BOB) would be replaced by ((LEG HUMAN) BOB), and this concept inherits from its generalizer, (LEG HUMAN), as desired. This replacement process is termed derivative subclassification. Note that if the data base contains

[ (NUMBER (LEG HUMAN)) #CHARACTERIZATION 2 ]

then by derivative subclassification (NUMBER (LEG BOB)) would be replaced by

((NUMBER (LEG HUMAN))((LEG HUMAN) BOB))

This takes place in two steps. First, since BOB is a HUMAN, (LEG BOB) becomes ((LEG HUMAN) BOB) by replacing LEG with (LEG HUMAN). Then, since ((LEG HUMAN) BOB) is a (LEG HUMAN), NUMBER is replaced by (NUMBER (LEG HUMAN)). The resulting concept inherits the characterization 2, as described.

A difficulty with this scheme is illustrated by considering not Bob's leg but Bob's left leg. If the latter becomes (((LEG LEFT) HUMAN) BOB) it is not in the class of (LEG HUMAN). If it becomes (((LEG HUMAN) BOB) LEFT) it is not in the class of ((LEG HUMAN) LEFT). One is required either to assert ((LEG LEFT) HUMAN) = ((LEG HUMAN) LEFT); to search the data structures, changing the order of the specializers; or to repeat specializers, (((LEG HUMAN) LEFT) HUMAN) BOB). Fahlman proposes to search, but to do parallel computation to avoid long search times.

Hayes (15) has suggested that one could avoid forming the concept (LEG BOB) and still find the number of legs BOB inherits from HUMAN.

We form

[((NUMBER LEG) HUMAN) #CHARACTERIZATION 2].

((NUMBER LEG) BOB) will then inherit from ((NUMBER LEG) HUMAN).

Hayes' method is in fact more awkward than indicated above because he identifies a context for Bob which is separate from Bob. Bob is represented both by a CONNIVER like context and a node in that CONTEXT. We avoid this by using concepts as contexts.

### 10. Semantic Significance of the Specializer

This section sets forth some basic conventions needed to establish a mapping between English words and phrases and OWL II concepts. Drawing on the discussion of syntagma in Section 3 one might propose that, for example, whenever two nouns like fire and plug are combined to form a compound noun, fire plug, then the corresponding concepts FIRE and PLUG should be combined by specialization to form the concept (PLUG FIRE). Unfortunately, compound nouns are often ambiguous, e.g. woman doctor or snake poison, while the definition of specialization stipulates that the genus and specializer of a concept be sufficient to uniquely specify it. This ambiguity can be dealt with in three ways.

- a) It can be ascribed to ambiguity in the constituents, e.g. river bank.
- b) The resulting phrase can be taken to idiomatically name other phrases, e.g. hot dog (skier), hot dog (sandwich).
- c) One can somehow distinguish "kinds" of specialization, e.g. woman doctor means either
  - i) doctor is characterized as a woman
  - ii) the doctor doctors women.

The first two can be added to any system containing the third.

In OWL II, "kinds" of specialization are realized by using meta-specializers. For example, instead of forming (DOCTOR WOMAN), DOCTOR is first specialized by a meta-specializer and then the result is specialized by WOMAN: ((DOCTOR CHARACTERIZATION.) WOMAN) or ((DOCTOR OBJECT.) WOMAN). The different meta-specializers, CHARACTERIZATION, and OBJECT, make possible disambiguation.

The eight types of meta-specializers are shown in Figure 8. These types have been chosen with several goals in mind.

<u>meta-specializer</u>	<u>example</u>	<u>abbreviation</u>	<u>use</u>	<u>English phrase</u>
INFLECTION.		*X	(DOG*X -S)	dogs
PREDICATE.		*P	(DOG*P FAT)	fat dog
APPOSITIVE.		*A	(DOG*A PET)	pet dog
SPECIES.		*S	(DOG*S BULL)	bull dog
INDIVIDUAL.		*I	(DOG*I "FIDO")	Fido
STEREOTYPE.		*T	(DOG*T LAP)	lap dog
SLOT.	OBJECT.	*SLOT.	(HIT*OBJECT. BALL)	hit ball
CONTEXT.		*C	(R23*C RADIO-1)	

Figure 8  
Meta-specializers

First, there are three principal ways that a concept C may be described by a concept B, using attachment.

- i) [ C \*PREDICATE B ]
- ii) [ C \*CHARACTERIZATION B ]  
[ C \*EXEMPLAR B ]
- iii) [ (OBJECT.\*ASPECT. C) \*CHARACTERIZATION B ]  
[ (LOCATION.\*ASPECT. C) \*PREDICATE B ]

C may be predicated by B, C may have B as a characterization or exemplar, or a concept representing a slot of C may be characterized or predicated as or by B. Attachment adds to the



description of an existing concept, it does not form a new concept. Frequently, however, one wants to form the concept corresponding to C described by B, e.g. one wants to further describe the class of tall men or pet skunks, or the procedure for hit baseball. To this end, the meta-specializers PREDICATE., APPOSITIVE., and the class of meta-specializers SLOT. have been created. Using these one can form:

- i)  $(C \circ P \ B) \equiv ((C \text{ PREDICATE.}) \ B)$
- ii)  $(C \circ A \ B) \equiv ((C \text{ APPOSITIVE.}) \ B)$
- iii)  $(C \circ \text{OBJECT.} \ B)$   
 $(C \circ \text{LOCATION.} \ B)$

The set of individuals described by  $(C \circ A \ B)$  is the intersection of those described by C and B.

In English grammar, it is traditional to distinguish between restrictive and non-restrictive modifiers. For example,

The philosophical greeks liked to talk.

can mean either that all greeks are philosophical and they like to talk (non-restrictive) or that only greeks who are philosophical like to talk (restrictive). The meta-specializers PREDICATE., APPOSITIVE., and SLOT. introduce restrictive modifiers.

It might seem appropriate to handle non-restrictive modifiers by attachment, but there is one complication which must be faced. This is the distinction between an expression and use of an expression raised earlier. When one says my friend, who is seventy, he means only for the description who is seventy to be attached to the description my friend in this particular use. Three ways this might be solved are:

- a) First specialize my friend by the current context before attachment.

- b) Introduce additional meta-specializers for non-restrictive readings.
- c) Attach who is seventy to another concept having the same extension as my friend in this use. That is, who is seventy is related to my friend only by a transient process and never shown in an explicit data structure (18, p. 62).

The choice seems to lead into discourse processing and won't be considered here.

The meta-specializers PREDICATE., APPOSITIVE., and SLOT. ascribe a precise semantic significance to the specializer. We know for example that any individual described by the concept (A\*P RED) must be RED. Thus, it is not necessary to also attach RED to this concept. In effect, when a concept corresponding to A described by B is formed, the semantic net link between A and B is replaced with an expression containing A and B. One has the option of forming new net links or new expressions.

A precise semantic significance can also be given to the meta-specializer CONTEXT., which stipulates that the specializer is the context of a concept. CONTEXT. gives the user the option of specifying context without giving any description of the relation of the concept to its context.

When a word like play is inflected with, for example, -ing, the resulting word playing is in the syntactic class of -ing, but takes its semantic properties from play. By giving -ING an INFLECTEE. slot, playing can be represented as (-ING\*INFLECTEE. PLAY). This is the appropriate representation for syntactic purposes since it is in the class of -ING. For semantic processing, however, it is useful to have a concept which is in the class of PLAY. This can be formed using the meta-specializer, INFLECTION., which is provided for this purpose. Playing can be represented either as (-ING\*INFLECTEE. PLAY) or as (PLAY\*X -ING).

The remaining meta-specializers, SPECIES., INDIVIDUAL., and STEREOTYPE., are all used with specializers of diminished semantic significance. How many people know, for example,



the relationship between skid row and skids or between bull dog and bulls. In these expressions the specializer is used primarily to identify the concept, not to describe it.

Notice that fat man is ambiguous between the sense of a man who is fat, and the idiomatic sense of a man who might work in a circus. The first sense may be pronounced with a slight pause between fat and man. The distinction can be seen in very fat man vs. circus fat man.

This distinction may be described as a choice between an adverbial and a lexical combination of the words. By definition, the adverbial relies primarily on the meaning of its constituent parts, while a lexical combination relies on the lexicon or memory for its interpretation. These two readings will be written (MAN<sub>o</sub>P FAT) and (MAN<sub>o</sub>T FAT).

The meta-specializer INDIVIDUAL. indicates that the concept describes an individual. SPECIES. is distinguished from STEREOTYPE. in that (A<sub>o</sub>S B) and (A<sub>o</sub>S C) are taken to describe mutually exclusive sets for any A, B, and C. (A<sub>o</sub>T B) and (A<sub>o</sub>T C) are not given this property. SPECIES. is useful in setting up a Linnean classification system.

It is not always easy to distinguish stereotypes from species, or, indeed, what should be stereotypes and species; but we have been able to make a practical distinction in the problems we have considered. Stereotypes always focus on one characteristic, e.g. he sits in the lap, he stays in the house, he barks. Species usually involve many attributes, e.g. dog vs. cat.

The mutually exclusive classification of species is done as a computational convenience. For example, one can quickly determine that an instance of a BULL-DOG is not a SHEEP-DOG because BULL-DOG and SHEEP-DOG are both species and neither is in the class of the other. On the otherhand, since FATHER is a stereotype it is not mutually exclusive with DOG and no such quick check is possible to determine that a DOG is not a FATHER. While one could



alternatively form individual specializations of DOG and characterize them with the stereotypes COLLIE, POODLE, etc. this would make it computationally more difficult to tell the breeds apart. The distinctions between breeds is perhaps not so important in general, but it would be, for example, in a computer program expert in the management of dog shows. The choice between species and stereotypes therefore depends in part on the particular expertise to be embodied in a given semantic model.

The usefulness of the species/stereotype distinction is based primarily on the computational capabilities of current computer systems. Since the computational capabilities of people differ from those of computers, it is difficult to say that the distinction is useful to people or even that they make it. The distinction does allow us to account for a phenomenon noted by Southworth (19).

"Similarly, a mutt is in one meaning a particular kind of dog (= mongrel), but in another meaning it is a way of talking about any dog (even a thoroughbred)."

We can form both meanings

(DOG\*S "MUTT")  
(DOG\*T "MUTT")

One might want to classify something into more than one set of mutually exclusive categories. For example, divide people by sex and also by occupation. To do this first stereotype people into people-with-a-given-sex and people-with-a-given-occupation and then form species of these two concepts.

It remains to note that snake poison can be either poison from snakes, or poison for snakes. The one falling in a class with snake skin the other with snake food. It has frequently been observed (20,21) that an English phrase such as poison for snakes or dog fights bulls can generally be found which picks the appropriate sense of the compound. Thus to disambiguate snake poison

one can assume that snake fills a slot of poison such as purpose or source. Alternatively, one can assume that snake specifies a predicate (OF-OR-PERTAINING-TO:OBJECT. SNAKE), which is then refined to appropriate senses before specializing POISON. The choice would take us into the representation of English in OWL II and will have to await a future paper.

## 11. Conclusion

The literature of artificial intelligence, linguistics, and philosophy contains many interesting ideas about the representation of knowledge. What is lacking is a proposal for a comprehensive system which is computationally appealing. In designing OWL II, we exposed many issues which any such system must confront. Those pertaining to the most basic level of representation have been presented here.

## 12. References

1. Lowell B. Hawkinson, "The Representation of Concepts in OWL", Proceedings IJCAI-75.
2. Peter Szolovits, Lowell B. Hawkinson, and W. A. Martin, "An Overview of OWL, A Language for Knowledge Representation," Presented at the Workshop on Natural Language for Interaction with Data Bases held by the International Institute for Applied Systems Analysis (IIASA) at Schloss Laxenburg, Austria, in January, 1977. Available as MIT Laboratory for Computer Science Report MIT/LCS/TM-86.
3. Randolph Quirk and Sidney Greenbaum, A Concise Grammar of Contemporary English, Harcourt Brace Jananovich, Inc. (1975)
4. H. Marchand The Categories and Types of Present-Day English Word-Formation, 2nd edition, Verlag C.H. Beck, Munich. (1969)
5. Daniel G. Bobrow and Terry Winograd, "An Overview of KRL, a Knowledge Representation Language," Cognitive Science, Vol. 1, No. 1, 1977.
6. Ira P. Goldstein and R. Bruce Roberts, "NUDGE, A Knowledge-Based Scheduling Program", IJCAI-77 p. 257-263.



7. Scott Fahlman, "A System for Representing and Using Real World Knowledge", MIT PhD Thesis in EE (Sept. 1977).
8. Brian C. Smith, "Levels, Layers, and Planes, The Framework of a System of Knowledge Representation Semantics", MIT EE MS Thesis (Jan. 1978).
9. William J. Long, A Program Writer, MIT Laboratory for Computer Science Report MIT/LCS/TR-187 (Nov. 1977).
10. G. J. Sussman, "Slices: At the Boundary between Analysis and Synthesis" to appear in IFIP WG 5.2 Working Conference on Artificial Intelligence and Pattern Recognition in Computer Aided Design, Grenoble, France, March 1978.
11. Marvin L. Minsky, "A Framework for Representing Knowledge", in The Psychology of Computer Vision, P. H. Winston, ed. McGraw-Hill, New York. (1977)
12. P. F. Strawson, Individuals, An Essay in Descriptive Metaphysics, Anchor Books (1963).
13. P. F. Strawson, "On Referring", Mind LIX, No. 235 (1950) p. 320-44.
14. Gerald J. Sussman and Drew V. McDermott, "From PLANNER to CONNIVER -- A genetic approach", Proc. FJCC 41 (1972) pp. 1171-1179 AFIPS Press, New Jersey.
15. Philip H. Hayes, "On Semantic Nets, Frames, and Associations" IJCAI-77 p. 99-107.
16. Richard Fikes and Gary Hendrix, "A Network-Based Knowledge Representation and its Natural Deduction System", IJCAI-77 p. 235-246.
17. John McCarthy, "Epistemological Problems of Artificial Intelligence" IJCAI-77 p. 1038-1044.
18. William Woods, "Whats in a Link", in Representation and Understanding, D.G. Bobrow and A. Collins, eds. Academic Press, New York.
19. F. Southworth, "A Model of Semantic Structure", Language, Vol. 43, No. 2, p. 342-361.
20. R. B. Lees, The Grammar of English Nominalizations, Indiana Research Center in Anthropology, Folklore and Linguistics 12, Bloomington, Indiana.
21. James R. Rhyne, "Lexical Rules and Structures in a Computer Model of Nominal Compounding", PhD Thesis in Computer Science, Univ. of Texas, Austin. (May 1976).
22. Robert C. Moore, "Reasoning about Knowledge and Action", M.I.T. PhD Thesis Proposal, Department of Electrical Engineering and Computer Science, 1977.
23. Lakoff, George, "Hedges: A Study in Meaning Criteria and the Logic of Fuzzy Concepts".



24. Brown, Gretchen P., "A Framework for Processing Dialogue", MIT Laboratory for Computer Science Report MIT-LCS-TR-182.