THE REPRESENTATION OF CONCEPTS IN OWL

Lowell Hawklnson

Project MAC
Massachusetts Institute of Technology
Cartridge, Massachusetts

Abstract

This paper discusses the theoretical basis of the formalist used as a foundation for OWL, a system for representing and processing conceptual knowledge. This foraallsa atteapts to capture the expressive power of natural language by adopting the underlying representational conventions and "conceptual models" of English. The foraallsa Is built around specialization, which is perceived to be the key organizing principle of English at a deep conceptual level. The use of specialization In combination with another low-level structural device, reference, leads to a slaple but powerful structure, the concept, which is ideal both for the representation of a broad spectrum of conceptual knowledge and for computation on existing machines equipped with large, high-speed, random-access aeaories.

Introduction

OWL Is a systea for representing and processing conceptual knowledge, intended for applications requiring interaction with huaans in natural language, at a human-like level of competence. Development of OWL has been proceeding for the past two years within the Automatic Programming Group of Project MAC under the direction of WIlliam A. Martin. It is his successor to an earlier systea called MAPL.

This paper discusses the theoretical basis of the <u>OWL formalism for conceptual knowledge</u>, which serves as a foundation for the OWL system. The purpose of this formalism is to allow conceptual (i.e., non-imagistic) knowledge to be expressed In terms of concepts and connections between concepts; in this respect, it does not differ froa such existing formalisms as Quilllan's seaantlc memory or Schank's Conceptual Dependency networks. It is quite different, however, in respect to details of structure, notation, and computer impleaentation. Among the advantages we expect to obtain froa the OWL foraallsa are:

(1) despite Its simplicity and uniformity, we expect it to be suitable for representing a broad spectrua of conceptual knowledge at all levels of abstraction and detail, in a "natural" way;

(2) it provides an organizational framework that should allow processing tasks required for comprehension of natural language to be carried out effectively and efficiently on a serial computer;

(3) its organizational fraaework is such that Its performance should not deteriorate significantly as the size of its knowledge base and the scope of its applicability increase, given an adequate amount of high-speed, random-access memory;

(4) the notation for it will be easy for English speakers to learn and use; and

(5) its structures and notation are especially effective for dealing with very large knowledge bases.

There are several distinct aspects of the OWL foraallsa which we Bight examine. This paper focuses primarily on the structural aspect, but also looks briefly at the notatlonal, English, and operational (i.e. implementation) aspects. For further description of the foraallsa, see Martin and Hawkinson.

To have any reasonable chance of realizing our goals for OWL in the near future, we have felt it wise to structure knowledge In OWL in accordance with the principles of organization of the human cognitive systea, insofar as we understand them, even though the continual advance of our understanding mandates continual evolution of the OWL formalism. Since auch of the significant evidence we have of this organization derives froa our understanding of the structure of natural language, we have tried to use those linguistic insights to develop a foraallsa that would capture the expressive power of natural language without sacrificing computational efficiency. We felt we could succeed in this only by adhering as closely as possible to the representational conventions and conceptual models of one particular natural language, English, even for the representation of "deep-level" structures, i.e., structures at the level of Schank's Conceptual Dependency networks or Minsky's frames. An obvious advantage of a foraallsa based on the structure of English is that translation back and forth between it and English might be done with relative ease; this is important not only for the OWL system as ultimately developed, but for the development process itself.

We do not yet know enough about the structure of English to have a definitive blueprint for our OWL foraallsa. We have had, in fact, to decide among competing alternative solutions to a host of representational probieas by applying principles of econoaical systea design — being careful, however, not to compromise our primary guideline of choosing representations that are readily mappable into English and vice versa. The structural organization we have arrived at as a result of our design efforts seems so simple and natural, and the evidence for Its pervasive use throughout natural language seems so intuitively convincing, that we feel it must hold profound implications for psychological models of human cognition. For example, it might be used as the basis for a strongly predictive model of word phrase association.

Though the OWL formalism has been implemented in LI5P, OWL does not in fact adopt or build upon

LISP's model of how data should be structured, namely, in terms of tree-like list structures and property lists of atonic symbols. Rather, the OWL formalism provides an alternative scheae for the low-level ("sub-frame") structuring of knowledge (hence data) that is, 1 believe, as fundanental in character as that provided by LISP. This formalism sight prove a suitable foundation not only for the OWL systen, but also for other high-level systeas like Moore and Newell's Merlin or any of the recently proposed "frame" systeas.

The Theoretical Basis for the OWL Foraalisn

In developing the OWL fornalisn for conceptual knowledge, the first big issue that confronted us was how to classify and "index" the enormous number of concepts OWL would need for huaan-like competence in natural language. We started with three fundanental principles:

(1) that essentially all conceptual knowledge should be represented in terns of a unifora, simple "building block", which we call a concept;

(2) that the set of all concepts should be arranged into a single conceptual taxonoay (concept tree) in such a way that individual concepts would "inherit" from superior concepts ■ost of their properties, especially "Beta-properties'* that indicate what to do with concepts in various interpretive contexts;

(3) that the vast bulk of the conceptual taxonoay should be deternined "derivatively", so that a few tens of thousands of individual classification decisions night suffice to determine a conceptual taxonoay containing millions of concepts.

Almost a corollary of the second principle is that the concept tree should have a low, relatively unlfora "fan-out" (where alaost all concepts have fewer than, say, ten iaaedlate subconcepts), since it is unreasonable to expect a concept to inherit aost of its properties fron concepts significantly ■ore general than it. A low, relatively unlfora fan-out also offers important structural and computational advantages.

We adopted these fundamental principles primarily because we Judged that a foraalism based upon then would simplify the task of building a large OWL knowledge base just enough to Bake that task feasible. However, we would not be surprised to find that these principles were also fundamental to the organization of human conceptual memory, which, after all, must acquire and maintain an even larger knowledge base.

Underlying the OWL formalism for conceptual knowledge, and hence the structure of a concept, are two basic structural devices: specialization and reference. Properly utilized, these devices permit the formalism to satisfy our three fundamental principles, and they also give It great expressive power and characteristics that make it efficient to work with. (To eliminate a possible source of confusion, note that when we speak of the structure of a concept, we understand it to be located at some one fixed place within a particular copy of a particular knowledge base, though it may be *referred* to in any number of places.)

Specialization

Specialization is a means for identifying a concept uniquely by a pair of entities: the concept's genus, a superior concept in the concept tree, and its speciallzer, aost often also a concept.* Every concept in OWL must, in fact, be identified by specialization, and thus a concept is often itself referred to as a specialization (of its genus). The significance, If any, of a concept's specializer cannot be determined according to a few simple, set rules; indeed, it may be arbitrary, though it typically depends on some generalization of the concept at or near the level of the genus.

The fact that every concept must have exactly one genus does not mean that a concept cannot be "a kind of" aore than one thing. Actually, as we shall later see, a concept can have any nuaber of distinct characterizations. This aeans that a genus could be viewed as simply a characterization singled out to provide identification for and primary classification of a concept. The choice of a genus Is nevertheless very Inportant, since in practice, aost low-level interpretive decisions as to what to do with a given concept Bust, for reasons of serial processing efficiency, depend on its genus alone. Fortunately, as we will demonstrate below for English, a natural language expresses many concepts in terns of an appropriate genus and specializer, and aost of the rest can be determined through the use of simple, productive, language-specific rules. Only a relatively snail nuaber of hard choices reaaln for the conceptual taxonomlst, but it is laportant that he Bake them wisely (though on philosophical grounds, it seems unlikely that there is one true conceptual taxonoay).

Specialization is directly evident in natural language; in fact, I believe it to be the aost important technique of signification (concept Identification) within natural language. Let us look at some English phrases that Identify concepts by means of specialization. In our examples below, we will underscore the part of each phrase that identifies the genus. The rest of each phrase, minus simple connectives like "of", "as", "than", and sometimes "to", "for", and "in", identifies the specializer (only because we have deliberately excluded extraneous modifiers, such as leading articles, that apply to the phrase as a whole). Of course, a genus or a specializer can itself be a specialization.

hit the ball, hitting the ball

get a book, get_ a Job, get to go,
get lost, get_ wet, get up

look up the name, put off my decision

*It has been observed that the term specializer can be misunderstood by persons who are Just becoming familiar with OWL terminology. It should be understood to aean "that coaponent of a concept which makes it special" or "that coaponent of a concept which makes it a particular specialization of the genus", aot as "that which produces the specializations of a concept".

go by bus, want him to know

flower garden, garden flower (Jespersen)

computer sales, August sales,
sales increase

sales of computers, sales for August,
increase in sales, sales in dollars,
big in size

glass of milk, can of beans, all of us

robins' nests, John's father,
John's being late

left hand, capital letter,
passenger tire

black tie (but not blue tie)

Friday, seller, wanting, wanted

sky blue, dead wrong, easy to read

big for a cat, bigger than an elephant,
less than 3

in the house, in June in trouble,
in toto, in stock, in good faith

in back of the bus,
in the back of the bus

let go of the block, let the block go

Let us now generalize from some of these sample phrases in terns of parts of speech, ignoring for now the possibility of exceptional cases.

verb direct object

verb infinitive, verb adjective,
verb particle (preposition)

verb secondary clause

noun nominal, nominal "of nominal,
genitive nominal

adjective infinitive, noun adjective

"as" adjective "as" nominal

preposition nominal, stem suffix

What criteria have we used here to recognize an English phrase as a specialization? First, the phrase must be meaningful as a concept, the test for which is that the phrase as a whole must denote something that could be further described. Secondly, the phrase must have a sub-phrase which one can say the whole phrase is "a kind of", though perhaps only in some very abstract sense (e.g., "In trouble" is a kind of "In" only in some abstract sense of the word "in"); the maximal such sub-phrase (e.g., "on top" rather than "on" in the phrase "on top of the table") identifies the genus of the phrase. Thirdly, the phrase must contain another sub-phrase which, when properly interpreted in context, combines with the genus to identify, independent of the context, the concept represented by the phrase as a whole; this other sub-phrase, If it exists, is the speciallzer. A couple of examples should help to clarify this last criterion. The phrase "the red book", though it satisfies the first two criteria, does not satisfy the third, since neither "the" nor "red" nor "the red" combines with genus "book" to unambiguously identify, independent of context, the book being referred to. The phrase "my father" on the other hand, would be a specialization since it has a genus, identified by "father", and a speciallzer, the referent of "my" (presumably determinable from context), which together identify the concept (though not necessarily *its* referent) uniquely without need of context. A useful diagnostic for English phrases that are specializations is that, in most cases, both the genus and speciallzer will carry stress, with the stress on the speciallzer *at least* as strong as that on the genus (when no word In the phrase is being stressed for the sake of emphasis).

In actuality, a large proportion of English grammatical constructions may be treated as particular forms of specialization, that Is, different syntactic patterns that all do basically the same thing, namely, identify a concept by specifying its genus and speciallzer. Analyzing English constructions as specializations, where possible, often obviates their analysis in terms of traditional grammatical categories, such as parts of speech. Where it is possible to analyze an English phrase as a specialization that "matches" a "sufficiently specific" concept in the knowledge base, there is no real need to also match It to some more abstract grammar rule.

Indeed, If we view the genus, speciallzer, and syntactic properties of each generic concept In an OWL knowledge base as a grammar rule, we can envision a conceptual grammar which, for a large knowledge base, could easily include hundreds of thousands of such rules. In a conceptual grammar, idioms would be rules, not exceptions to rules. A conceptual grammar acknowledges and emphasizes the essentially idiomatic character, at every level, of natural language as used, whereas traditional grammars have tended more to emphasize the essential regularities that can be abstracted from instances of its use. From our point of view, in fact, the principal goal of a traditional grammarian is to approximate a large conceptual grammar by a grammar of relatively few rules expressed in some particular formalism, a task of immense difficulty. Traditional grammars, even If they could be realized, would not be as satisfactory for OWL as a more straightforward conceptual grammar. Because most of what is needed for a conceptual grammar is already required to be in the knowledge base for purposes of conceptual modeling, it would almost certainly be less work to create a conceptual grammar for a large knowledge base than to program the mapping between OWL conceptual structures and structures produced by a parser for some traditional grammar.

Thus far, we have analysed as particular forms of specialization only those English words and phrases wherein both genus and speciallzer are manifest. However, for the sake of uniformity, we would like to treat monomorphemic words as specializations and also many specific phrases and constructions that fail to express either the genus or the speciallzer or both. Thus, we typically treat the concept corresponding to a

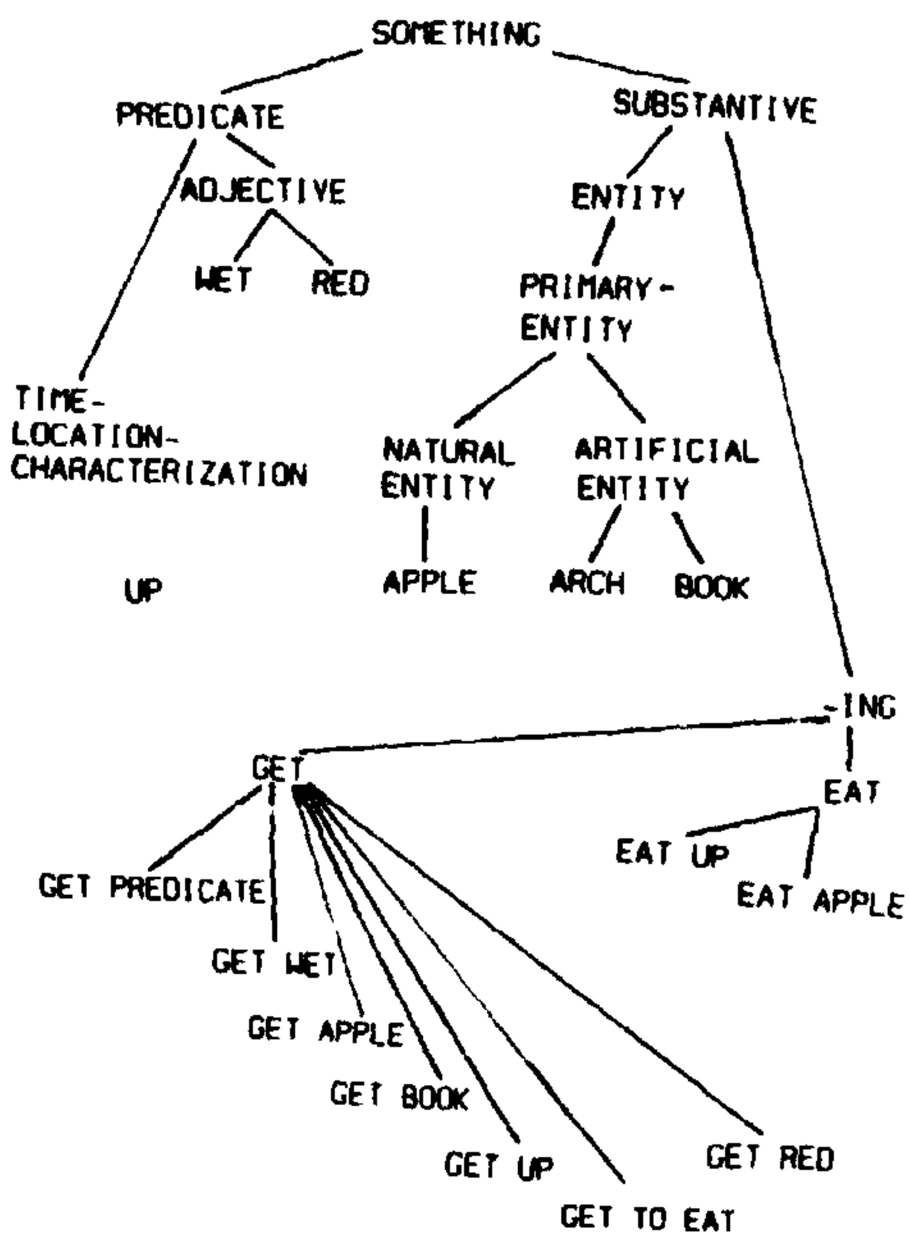

FIGURE 1. A Tiny Concept Tree

monomorpheaic wo*,d as some appropriate genus specialized by the written syabol for that word, e.g., the concept spider as the concept insect specialized by "spider." Non-word aorpheaes, such as "-ing\ are treated siallarly. We call such concepts <u>syabol aeanIngs</u>. (Concepts specialized by concepts, on the other hand, are said to be <u>classified:</u> in this case, the speclaiizer Bay be referred to as a <u>classifier</u>.) An idlomatlc concept like "hot dog" is dealt with in a way that is soaewhat analogous to the treatment of syabol aeanIngs; in particular, "hot dog" is treated as the concept sandwich specialized by the concept consisting of concept dog specialized by concept hot. (Note that there is another teaning of "hot dog" that would require the concept skier in place of the concept sandwich.) More generally, we have observed a dlachronic tendency for concepts to becoae identified in English by their specializes, e.g. "capital letter" by "capital" and "New York City" by "New York", where both forms are in common use, and "general officer" by "general", where the first fora (the one revealing the genus) is now archaic. As Martin points out, this tendency to identify a concept by its speclaiizer is just a special case of the widespread, aore general phenoaenon In English of a concept being identified (naaed) by some other, usually "closely related" concept.

We could continue to exhibit here successively aore elaborate foras of specialization, but that would take us beyond the scope of this paper and into Martin's paper, which presents a far aore comprehenslve theory of English. As It Is, our treataent above of "spider" and "hot dog" cannot be argued here to be anything aore than conventional; a deeper Justification could not possibly be given without that aore coaprehenslve theory.

We have defined and discussed specialization as a Beans for *identifying* concepts, but specialization Bay be equally viewed as a Beans for *classifying* concepts. Both the genus and speclaiizer of a concept contribute to that concept's classification. The contribution of the genus is readily apparent- the gross structure of the concept tree can, in fact, be defined as the set of ail mappings from concepts to their genuses (or, alternatively, froa concepts to their specializations). This gross structure is then further refined by Baking use of the specializers, a process known as derivative subclasslflcatlon. The rule for derivative subclasslflcatlon of concepts by their specializers Bay be stated, soaewhat iaprecisely, as follows: subclassify the specializations of a concept as their specializers are classified.

An illustration will serve to clarify the aeaning of derivative subclassification. Figure 1 shows the gross structure of a tiny concept tree (or a tiny extract froa a large concept tree -- the analysis applies in both cases). Figure 2 shows that saae structure refined by derivative subclassification, i.e., the fine structure of the tree. Notice that new generic concepts have been Introduced into the tree such that, for every pair of classified specializations that share a coaaon genus, say g specialized by s, and s_2 there will exist a corresponding specialization of g by the aost specific coaaon superior of s_1 and s_2 except when that superior is the suaaua genus "somethlng". Thus, for instance, concepts "get adjective" and "get substantive" are introduced,

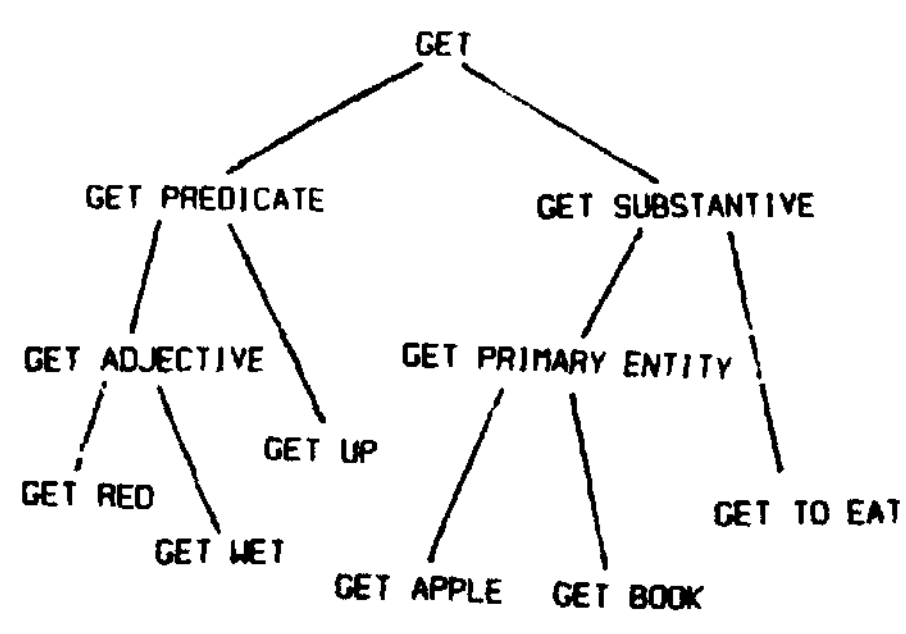

FIGURE 2. The Tiny Concept Tree of Figure 1 Refined by Derivative Subclassification

but "get entity" and "get natural-entity" are left out (although one or both could later be introduced either directly or in consequence of other newly formed specializations of "get").

To facilitate further discussion of the fine structure of a derivatively subclasslfled concept tree, we need to define a few additional terms. To begin with, we will define the term generaUzer to be the immediate superior of a concept in a derivatively subclasslfled concept tree. Thus, the general izer of "get book" in the concept tree of Figure 2 would be "get primary-entlty". Note that, unlike the genus, the generalizer of a concept may change as new concepts are introduced into the concept tree. The generalization path of a concept is the sequence of concepts encountered in moving from it to the SUBBUB genus by successively taking the generalizer at each step. Again with reference to Figure 2, the generalization path of "get wet" would consist of "get adjective", "get predicate", "get", "-ing", "substantive", and finally "something". The generalization path of a concept is the primary path along which properties are inherited; we would therefore expect "get wet", for instance, to Inherit certain of the properties of "get adjective", "get predicate", etc. (Some properties are not inherited by inferior concepts because they are contradicted at lower levels or because they are specifically known not to apply beyond a certain depth). We will use the tem conceptual aodel to refer to a concept and all its subconcepts. Thus, in the concept tree of Figure 2, there is a conceptual aodel for "get", which consists of "get" and the subtree beneath it, together with whatever night be attached thereto. Finally, generic subconcepts in a conceptual model are said to be senses of the concept at the top of the nodel. Hence, "get adjective" and "get substance" could be described as senses of "get".

What is the significance of using a derivatively subclasslfled concept tree? First of all, if we were to have a conceptual taxonomy with tens of Billions of concepts (which we would surely need to approach a hunan level of competence in natural language), we might expect derivative subclasslficatlon to account for at least 99% of the classification decisions required to build that taxonoay. Secondly, if, when building our taxonomy, we contrive to put a limlt (say ten) on the number of *symbol meaning* specializations a concept may have, then that saae lialt can be shown to hold for the nuaber of immediate *classified* specializations any concept could possibly have. Thus, we can easily arrange to get a concept tree whose fan-out at any given node is strictly Halted. Thirdly, If we put into our concept tree a representative sample of the specializations of a particular concept (as they occur in ordinary use of the language), derivative subclasslflcatlon will provide us with a useful and intuitively reasonable set of senses of that concept to attach inheritable properties to. Thus, if our saaple of the specializations of "have" includes, say, "have a tail", "have four legs", and "have a top", where "tall", "legs", and "top", are already classified as "parts", then derivative subclasslflcatlon will give us the sense "have as part", which, as the head of an important conceptual submodel of "have", will carry many inheritable properties. (The question of how we might *automatically* derive properties for "have as part" from its inferior concepts is interesting, but beyond the scope of this paper.) Note that just because a concept like "have a tail" is a subconcept of "have as part", there is no need to explicitly spell out "have as part" when expressing "have a tail" in a formal notation, any more than there is a need to do so when expressing it in English. In fact, we might go so far as to say that a rule of derivative subclasslflcatlon is necessary in a very large, finely subdivided taxonomy to allow economical identification (naming) of concepts within it. Finally and in sua, derivative subclasslflcatlon furnishes the Beans for satisfying the three fundamental principles proposed near the beginning of this paper.

Reference

Reference, the second of the two basic structural devices of OWL, is the Beans by which one concept or syBbol aay refer to another. By the ten reference, we are trying to suggest the general function of a connection that is syabolic (non-physical), unidirectional, precisely definable, and differentiated only as to function (not, say, by a label on the connection); these iaportant properties of reference are not conveyed well by such tens as "pointer", "association", or "associative link". The Bost important computational property of references is that the references under a concept or symbol (to other concepts or syabols) aay be accessed trivially froa it. In fact, a concept aay be thought of as being coaposed entirely of references, arranged In an orderly way. (A syabol has a spelling in addition to references).

We distinguish three categories of reference: the generalizer and speclallzer of a concept, Indexing, and description. Except for generalizer and speclallzer references, which are recognized by where they appear in a concept's structure, the category of a reference Bay be determined solely by inspection of the referent (the concept referred to). Of course, in an actual computer implementation, we might use a more elaborate structure for a concept than is strictly necessary, so that we can categorize a reference or retrieve all references of a particular category aore rapidly than by inspection of the referent. It is well to be aware, though, that such elaborations of structure are not theoretically activated or required.

indexing in OWL is a category of reference used to connect (1) any concept to Its immedlate specializations, (2) any concept to all concepts specialized by it, or (3) any syabol to its syabol meanings. Thus, with indexing, it is easy (coBputationally) to go froB the syabol "spider" to the concept spider, and froa there to specializations like red spider and black widow spider, and also to concepts classified (specialized) by spider, such as spider web, spider monkey, and spidery. An index reference can be recognized as such by the fact that the concept or symbol It is under will appear as either the generalizer or the speclallzer of the referent. Thus, a reference to spider web under spider would be identified as an index reference by the fact that spider is the speclallzer of spider web. Note that index references under a concept c are sometimes called "back pointers" because they "point back at" concepts whose generalizer or specializer "points at" c.

There is one unresolved problem with index references of type (2) above. Whereas index references of types (1) and (3) are Halted in number by derivative subclasslflcatlon and by the low level of morphemic ambigulty in English, respectively, there will exist in the knowledge base nany concepts having hundreds of index references of type (2), e.g., the concept house. This goes against a principle we would like to be able to adhere to, naaely, that the number of references under any one concept should never significantly exceed,, say, twenty. The likely solution to this problem is a second form of derivative subclasslflcatlon to subclasslfy all the concepts specialized by a given concept. However, we have not yet adopted such a solution, in part because it turns out to be unsatisfactory to use the obvious counterpart of the rule for derivative subclasslficatlon of the concept tree.

Pescription is a broad category of reference that, by definition, includes every reference not used for indexing or for the generalizer or speclalizer of a concept. The tera description is nonetheless quite appropriate, since individual descriptive references (or descriptors) can alaost always be said to describe the concept they're under (the subject of the description). We will distinguish here, on functional grounds, two principal types of descriptive reference, naaely characterizations and modiflers, each with nuaerous subtypes. But we will side-step, for now, the difficult problea of providing effective criteria (tests) for such distinctions, relying instead on illustrative examples and on the noraal English aeanings of the terms we have chosen for the types And subtypes.

A characterization is an alternative, partial description of a concept (the subject of the characterization). For example, a particular dog might be characterized as follows: a nuisance (abstract characterization), Mary's pet and Fido's father (relational characterizations), a good swimmaer (skill characterization), a chaser of cars (habitual role), and the dog who ran across our yard yesterday (event participant).

Relational characterizations deserve special note. In OWL, a relationship a A a is represented by a two-way characterization: a as a value characterization under *R* specialized by a, and R specialized by *b* as a relational characterization under a. Thus, for instance, "Ellen is the mother of Sam" would be represented by the value characterization "Ellen" under "Sam's mother" In conjunction with the relational characterization "Sam's mother" under "Ellen". Seaantic case relationships such as "New York was the location of the deaonstration" and arlthmetic relationships such as "5 is greater than 3" would be handled in an analogous way. This composlte technique for representing relationships, which has been borrowed directly froa English, offers at least two important advantages over the traditional logical formulatlon widely used as the basic connective link in seaantlc networks. First, specialization of the relation by one of its arguments produces a conceptual model for the relation; such a conceptual aodel serves as an effective fraaework for organizing knowledge of specialized uses of the relation. For exaaple, the conceptual aodel of "father" Bight be expected to include the senses "father of a person", "father of a country", "father of a science", and "father of a bride", each appropriately described. Second, whenever there are relationships that are identical except for their value arguments, the set of distinct value arguments will be automatically grouped as value characterizations under the shared relational characterization. Thus, if Mary has both a dog and a cat, they would both appear as value characterizations under "Mary's pets". We might also point out here that OWL does not represent reversible relationships canonically, e.g., "the elephant is bigger than the mouse" would be represented differently froa "the aouse is smaller than the elephant"; humans also seem not to canonicalize such relationships.

A aodlfler is any descriptor not deemed to be a characterization of its subject. Typical aodiflers are: "black" (property), "actual" (feature), "the" (determiner), "all" (quantifier), and "in the box" (location). Note that sometlaes the dividing line between modifier subtypes is not sharp (e.g., between properties and features).

A descriptor that is often applied to a particular concept tends to become a speclalizer of that concept. For example, "fat man" is specialized by a modifier ("fat") and "father figure" by a characterization ("father"). A concept specialized by a descriptor is usually more specialized in aeaning than if that descriptor was acting *only* as a descriptor; thus, not every black bird is a blackbird, nor is every round house a roundhouse.

An laportant transformation in OWL, again a reflection of a similar transformaton in English, is the conceptualization of a description to produce a predication. (We use the tera predication here in the classical sense of a predicate applied to, i.e., specialized by, some subject). A description (descriptor under some subject) Is transformed into a predication as follows: first, the predicate Is formed as a specialization of "being" (the copula) by the descriptor; then, that predicate is specialized by the descriptor's subject to give the predication. The reason for conceptualizing (making a concept out of) a description is so that the description can itself be described, specialized, or otherwise used as a concept. In linguistic terminology, both the process of conceptualization and the resulting predication would be referred to as noainalization. An exaaple of a predication in English is the phrase "John's being happy", which we have already used above as an exaaple of speciallzation.

The Structure of the Knowledge Base

All knowledge in an OWL-based system is held in a single, large, unified knowledge base of concepts and syabols, except for a rather saall though laportant amount embedded in LISP and machine language programs and their associated data structures. This knowledge base is not structurally partitioned by subject matter, by permanence (long-term vs. "interaediate-tera"), by type of concept (seaantlc vs. episodic, individual vs. generic), or by level of abstraction (surface language vs. primitive action, for exaaple), although obviously such distinctions aust often be inferrable by inspection of individual concepts. There is even no separate lexicon -- knowledge about English words and grammar Is distributed throughout the knowledge base and is essential to its organization.

All of the following terms are virtual synonyms of knowledge base: world model, semantic network, Knowledge Net (a Merlin term), conceptual data base, conceptual memory, conceptual grammar, conceptual taxonomy, and concept tree. The terms conceptual taxonomy and concept tree do not encompass symbols, but then symbols play only a minor role in the knowledge base. Symbols correspond to English words and morphemes and contribute no more to the knowledge base than do the spellings of words to the content of a book. Essentially the entire body of knowledge is represented in terms of concepts.

When OWL structures are modeled in terms of L15P data structures, a concept is typically represented by a list of its constituent references, where the first two list elements are the generallzer and speclalizer, respectively. A symbol is represented by the atomic symbol whose "pname" shares its spelling, and the references under the symbol appear on a list carried as the value of a property of the atomic symbol.

Relation to Other Work

Various uses of specialization, especially for compounds, have long been recognized by grammarians. Bloomfield, for example, discusses how the meanings of various kinds of compounds depend on the meanings of their constituent parts, thereby illustrating many of the properties of specialization. (He also compares his classes of compounds to similar classes used by Sanskrit scholars over two thousand years ago). Jespersen saw the "specializing power" of a modifying phrase (one of the lower "rank"), noting that "the object serves to make the meaning of a verb more special". However, no one appears to have understood the universality of specialization as a means for identifying and classifying concepts expressible in natural language, even though binomial systems of nomenclature have been used to identify and classify biological species ever since Linnaeus introduced the first such system in 1753.

Conceptual taxonomies ("hierarchies of knowledge") are of classical origin. Raphael was perhaps the first person to consider using a conceptual taxonomy as part of a computer system interacting with a user in English. He rejected it, however, on pragmatic grounds, citing the complexity of the required structure and the difficulty of producing a "useful" taxonomy. During that same period, Quilllan made a serious attempt to design a model suitable for a large semantic memory, but his organizing principles were too weak and his structure was connected in too ad hoc a fashion to be effective in dealing with knowledge on a large scale. There has only recently been a revival of interest in using large conceptual hierarchies, triggered primarily by disenchantment with the poor efficiency and non-Intultive functioning of systems using logic-based rules of inference. (Inheritance of properties, which is primarily what hierarchies are good for, is probably the simplest and most "natural" rule of inference). However, the hierarchical structures proposed by Moore and Newell, winograd, and Fahiman pernit concepts to belong to any number of classes, none of which is particularly favored, whereas in OWL each concept is considered to belong to only one class (its genus), though it may have any number of characterizations (each itself a distinct concept). The OWL use of hierarchy is superior if a genus can be readily chosen for each concept and if inheritance of properties along the generalization path of a concept is often sufficient for making low-level decision as to what to do with that concept.

To the best of my knowledge, the closest relative of the OWL formalism for conceptual knowledge is found in the Merlin system of Moore and Newell. As a prelude to making comparisons between these two formalisms, let us tabulate corresponding terms:

OWL	Merlin
knowledge base	Knowledge Net
conceptual taxonomy	hierarchy of knowledge
concept	p-structure
genus	schena
description	component
characterization	alternate view

Note first that in Merlin there Is no equivalent of specialization as an essential aspect of every concept in the knowledge base. Thus Merlin has no universal means of concept Identification (aside from full specification of the schema and all components) and no fine structuring of the concept tree through derivative subclasslflcatlon. (Actually, genus-speciallzer identification of concepts can be found among samples of Merlin fi-structures, e.g., [EYES BLUE] and [TASK LOGIC), although other p-structures of a similar form, e.g., [AUTHOR NEWELL-SIMON-SHAW] and [+ 3], do not follow the same [genus speclalizer] paradigm, at least according to OWL's English-based criteria for what constitutes a specialization). Second, each alternate view of a p-structure can adequately represent It as a principal view, whereas a characterization of an OWL concept is an entirely separate concept that need be only a partial description and which could be the characterization of any number of other concepts. Mapping among alternative characterizations is not a fundamental operation in OWL as it is In Merlin. Thirdly, OWL represents essentially all knowledge as concepts in a single, unified knowledge base, whereas Merlin uses a separate formalism for expressing procedures (actions), which are, however, attached to and considered part of the Knowledge Net. Neither OWL nor Merlin makes a basic structural distinction between individuals and generics.

Acknowledgements

! would like to thank ay colleagues Gretchen Brown, Suzanne Hawkinson, William Long, and Alex Sunguroff for their valuable suggestions and criticism, and especially Williaa A. Martin, whose work In the area of natural language originally led to and then later extended the theory reported herein.

This work has been supported in part by the Defense Advanced Research Projects Agency work order 2095, under Navy Contract IN0014-75-C-0661.

Bibliography

1. Bloomfield, L., *Language,* Holt, RInehart and Winston, New York, 1933, pp. 207-246.

2. Fahlman, S., "Thesis Progress Report: A System for Representing and Using Real-World Knowledge**, Meao No. 331, Artificial Intelligence Laboratory, M. I.T., Caabridge, Mass., 1975.

3. Hawkinson, L., "The OWL Foraalism for Conceptual Knowledge", Technical Memorandum, Project MAC, M.I.T., Cambridge, Mass., 1975.

4. Jesperson, 0., *Analytic Syntax,* Holt, RInehart and Winston, New York, 1969, p. 124.

5. Martin, W.A., "Conceptual Grammar", Autoaatic Programming Group Internal Meao 20, Project MAC, M. I.T., Caabridge, Mass. (submitted to the fall 1975 meeting of the Association for Computatlonal Linguistics).

6. Minsky, M., "A Framework for Representing Knowledge", Meao No. 306, Artificial Intelligence Laboratory, M. I.T., Caabridge, Mass., 1974.

7. Moore, J. and Newell, A., "How can Merlin Understand?", In Gregg, L. (ed.), *Knowledge and Cognition,* Laurence Erlbaua Associates, Potomac, Md., 1973.

8. Quillian, R., "Semantic Memory", in Minsky, M. (ed.), Semantic *Information Processing,* M.I.T. Press, Cambridge, Mass., 1968.

9. Raphael, B., "SIR: Semantic Information Retrieval", in Minsky, M. (ed.), Semantic *Information Processing,* M. I.T. Press, Cambridge, Mass., 1968.

10. Schank, R., "Identification of Conceptualizations Underlying Natural Language", in Schank, R. and Colby, K. (eds.), Computer *Hotels of Thought and Language,* W. H. Freeman, San Francisco, 1973.

11. WInograd, T., "Fraae Representations and the Declarative/Procedural Controversy", in Bobrow, D. and Collins, A. (eds.), *Representation* and *Understanding,* Academic Press, New York, 1975.