

STRUCTURED INHERITANCE NETWORKS AND NATURAL LANGUAGE UNDERSTANDING

Henry H. Leitner*
Center for Research in Computing Technology
Aiken Computation Laboratory
Harvard University
Cambridge, MA 02138

Michael W. Freeman
Advanced Development Organization
Federal and Special Systems Group
Burroughs Corporation, P.O. Box 517
Paoli, PA 19301

One of the most salient difficulties facing the designer of a Natural Language Interface is the appropriate characterization of the semantics and pragmatics of new domains. Whereas various syntactic parsers have been developed which are largely domain-independent (at least in principle), the possibility of accomplishing something similar in the realm of semantics and pragmatics has never been convincingly demonstrated. Focusing on the various epistemological bases for associating attributes with entities, we examine some of the requirements that this suggests need be met by a system which claims to provide a domain-independent conceptual schema core, from which domain-specific extensions can be interactively elaborated. In addition to showing how our approach differs from one which tries to arrive at canonical representations in terms of invariant semantic primitives, we sketch out ways in which it can be integrated into a general parsing mechanism.

0. INTRODUCTION

An ultimate goal of researchers in natural language (NL) processing is to have a system that interacts appropriately with human users in their own natural language. However,

- it is unfeasible at present to construct a general system which can handle all NL inputs in every domain and situational context.

Thus a revised goal might be a system that interacts intelligently with certain types of users in their own NL in particular universes of discourse (e.g., PARRY, LUNAR, SOPHIE, PLANES, SAM, ...). However,

- users may find it difficult (if not impossible) to make themselves understood in various instances;
- the effort involved in converting such systems to new domain and context coverage is considerable.

For these reasons, a revised goal might be a system that can be easily extended or customized by users, at run time (e.g., LIFER or REL). However,

- to be generally effective, this may demand that a user be more linguistically and computationally sophisticated than can be reasonably expected.

Thus a revised goal might be a system that can be easily shifted to new domains by its designers or maintainers (e.g., ROBOT). However,

- the real expertise concerning a new domain is generally not in the hands of the NLI (NL interface) designers or maintainers.

Thus a revised goal might be a system that can be conveniently shifted to new domains by "knowledge base experts" who need not also be linguistic or computational specialists (e.g., KRL or KLONE). However,

- there is generally the need to start over each time from scratch in order to avoid introducing inconsistencies and incomprehensible conceptualizations into the pre-existing knowledge network.

Thus a fifth goal revision is to have a system which can be easily adapted to new domains by knowledge base experts who work within a conceptual framework that builds upon a core of very abstract concepts which can be further generalized or specialized to represent domain-specific concepts, viewpoints, etc. This raises a number of questions that we will address in the course of this paper:

- Is it possible to develop a "core" knowledge network (KNET) that can fulfill the function proposed for it above?
- Assuming that we can develop such a KNET, what additional proficiency would a domain-expert or an interactive system need in order to expand it appropriately for a particular domain or application?

* This work was supported in part by the Rome Air Development Center under contract F-30602-77-C-0197.

- (diamond shape) which expresses interrelationships between the conceptual subparts. As described by Brachman, "these relationships [i.e., the role descriptions plus structural descriptions] give the Concept its 'gestalt'."

Whenever two concepts are in a generalization/specialization relationship to one another, the features they have in common can be factored out in such a way that only the differences need be represented. This is illustrated in Figure 2, where the AKO (a-kind-of) link represents a channel along which common properties can be inherited either intact, restricted or differentiated from a more general to a more specialized concept. Whenever a concept is specialized to the extent that all of its role-description nodes have values associated with them, it takes on a special status as an individuated concept. This is denoted by a shaded oval (cf. "HARVARD" in Figure 2), and is connected to its parent concept via an INSTS link.

1. OVERVIEW OF A STRUCTURED INHERITANCE NETWORK FORMALISM

Brachman developed a knowledge representation notation (SI-Nets) which has a very clear and explicit epistemology. According to Brachman, "Structured conceptual objects are organized in lattice-like networks, with inheritance between descriptions carried by structured 'cables'."* An example conceptual structure is illustrated in Figure 1:




FIGURE 1: SIMPLE SI-NET

The oval shapes denote conceptual entities, while the role nodes (squares) are intended to represent the intensional description of potential role fillers in context. The third component part of the formalism is the structural description




FIGURE 2: THE AKO LINK

It should be stressed that Brachman's primary goal was to make explicit the semantics of a representation system, thereby overcoming the expressive inadequacies of the many "semantic network" node-arc schemes in recent years. He makes no claims whatsoever about the semantics of what is being represented -- any concept can be represented in the SI-Net formalism any way the user chooses. Thus a particular concept might be defined in functional terms or in terms of its major parts and subparts, etc. The semantics of the world one thinks [s]he is modelling are not implicitly embedded in Brachman's system -- but the semantics of the representation system itself are clear and precise.

* Brachman, Ron: "Theoretical Studies in Natural Language Understanding." BBN Report #3888, Cambridge, Mass., September, 1978. p. 25.

2. EXTENSIONS TO THE BASIC NETWORK FORMALISM

The preceding remarks raise some rather basic issues, such as: what does it mean for a concept to have an attribute or to have a role; how do the latter become associated with concepts originally, etc. One reason we need to concern ourselves with such issues is that if a taxonomy of the ways in which different types of attributes can be associated with concepts is developed, then there is hope -- at least in this area -- of being able to give a knowledge base administrator (KBA) automated assistance in creating or extending a core semantic network. This matter is discussed further in Section 4.

Our point of departure here is the observation that the initial association of attributes with entities results from the creation of something which did not exist as such before. This does not necessarily entail the creation of a new physical object. In the context of human society, for instance, one is usually concerned with the creation of new social, political or legal entities. Thus, being born into a particular society transforms a mere (human) being into a societal entity having attributes such as NAME, RESIDENCE, MARITAL-STATUS, and so on. Note that we are not speaking of the association of a particular value with an attribute, but rather of the attribute itself with a new entity, within the framework of a highly developed social or legal order.

In this framework, the relationship of states resulting from specific events takes on fundamental prominence. For example, as the result of an instance of the "HIRE" event concept represented in Figure 1, a state of "EMPLOYMENT" comes into being. The semantics of the DATTR/RESULT arrow in Figure 3 is meant to capture this relationship.




FIGURE 3: THE DATTR-RESULT LINK

In turn, this allows us to conceptualize an EMPLOYEE as a societal person put into a state of employment as a result of playing the role of the "second party" (or q-agent) in a particular type of contractual event, namely that represented by the "HIRE" concept in Figure 1.* The QUA link in Figure 4 is the means by which we represent this interdependence of relationships.




FIGURE 4: THE QUA AND VIZ LINKS

The QUA link is the primary channel over which new attributes come to be attached to already existing entity concepts. Note that this is not simply a matter of viewing a given entity from some particular perspective, although it does give rise to such possibilities, as illustrated by the VIZ links in Figure 4. In this latter case, however, the perspective is taken as already having been intrinsically established, whereas in the former case what is being represented is how such perspectives are actually to be defined. This distinction between VIZ and QUA links becomes quite clear when one addresses the problem of making explicit the semantics of inheritance associated with each of the two link types.

The VIZ link has precisely the same inheritance properties as the AKO link. The only difference is that it serves to partition the multi-dimensional attribute space of a given entity type into proper subsets. This allows one to distinguish higher level entity concepts in the "gestalt" of a lower level entity concept having multiple inheritance paths. The reason for doing so, however, is contained within the semantics of the QUA link.

Since QUA links can derive from role description nodes within verbal concepts, the relationships between the parent role node and other role nodes within the verbal concept open up new paths for attributive inheritance that are not present in just the AKO (or VIZ) links. In the case of EMPLOYEE, for example, as illustrated in Figure 4, attributes such as EMPLOYER, DATE/PLACE-OF-HIRE, etc., result not from its AKO relationship to

* The importance of modelling ROLE-ENTITY associations has also been stressed by Bachman and Daya: "The Role Concept in Data Models." Vldb-3, Tokyo, Japan, October 6-8, 1977.

SOCIETAL-PERSON, but rather to the other role description nodes in the TO/BE/EMPLOYED concept, which themselves derive from the TO/HIRE concept via the DATTR/RESULT link illustrated in Figure 3. Now, by descending into the structural-description (S/D) nodes of these verbal concepts, one arrives at conceptual representations of how the various roles subsumed under the verbal concept are themselves related to one another. Wherever an S/D component co-references the (q-agent) node in TO/HIRE to which the EMPLOYEE concept is ultimately rooted, one has a potential source for generating a VIZ link. If one regards a TO/HIRE event, for example, as consisting of a job-offer and a job-acceptance, in which the applicant (q-agent) commits himself or herself to perform certain services in return for certain recompenses over some sort of time-span, then for each service to be performed, the employee can be viewed as its agent (e.g., WORKER), and for each recompense to be received, the employee can be viewed as its benefactive (e.g., PAYEE, GROUP-INSUREE, etc.). These are precisely the perspectives that we have represented via the VIZ links in Figure 4. In this way, an attribute such as SALARY comes down the VIZ link from PAYEE to EMPLOYEE, just as JOB-SITE does from WORKER, whereas RESIDENCE comes down the AKO link from SOCIETAL-PERSON, and EMPLOYER over the QUA link from TO/BE/EMPLOYED via the DATTR/RESULT link from TO/HIRE.

The discussion up to this point has focused on how attributes come to be associated with entities through performance events which result in certain socially or legally recognized states. We have said relatively little about the specification of domains from which these attributes can take on meaningful values. We have simply placed a value/restriction (V/R) link on role-description nodes (pointing to such concepts as SOCIETAL-PERSON, LEGAL/ENTITY, PERFORMANCE), and have relied on the indulgence of our audience to accept this as obvious. This has led us to gloss over the determination of the AKO-link between a QUA-derived concept and the V/R concept associated with the role-description node to which the QUA points. Yet it is not necessarily obvious in what sense an essentially functionally specified entity can be regarded as "a-kind-of" object pure and simple. Put another way, we need to look more closely at the implications of saying that the event of casting an entity into a particular role results in the "creation" of a new entity through its association with attributes (not just attribute-values) which it did not previously possess. In what follows, we will briefly consider how this approach leads one to view the domain specification (V/R) of attributes of this type as an event-driven state-transition network,

from which secondary attributes can be derived, as dictated by whatever the pragmatic needs of a given user community might be. This also provides an extremely straightforward means for specifying certain types of data base integrity constraints, update histories, identity conditions on transformable objects, etc.

In Figure 3 we represented the concept TO/BE/EMPLOYED as connected to the event concept TO/HIRE via a DATTR/RESULT link. We're all aware, however, that one remains employed by a particular legal entity only so long as certain other events do not occur, e.g., being fired, laid off, released, retired or dying. Whenever any of these other events occurs, one enters a corresponding new state. In fact, the history of this progression from one state to another via events of the types just mentioned can itself be viewed as the value of some higher level attribute, say, EMPLOYMENT/STATUS. Just as a SOCIETAL-PERSON acquires an EMPLOYER attribute through becoming an EMPLOYEE (cf. Figure 4), so a new-born human acquires an EMPLOYMENT/STATUS through becoming a SOCIETAL-PERSON as a result of being born into a particular type of society. Part of the "definition" of such a society is precisely the types of attributes that are associated with individuals upon their birth into it. And the domain-specifications of most of these attributes are most naturally represented as constellations of event-mediated state-transition networks, such as that illustrated in Figure 5 for EMPLOYMENT-STATUS.




FIGURE 5: STATE AND EVENT MONITORS

For any instance of a SOCIETAL-PERSON, the actual value of its EMPLOYMENT/STATUS is the current state it happens to be in. Since the latter results from a specific event-instance, one could just as well regard the event itself as the primary value, from which a current state value can always be computed on demand. The converse, of course, does not hold, thus underscoring the derived or secondary nature of state attributes such as EMPLOYMENT-STATUS. Let us therefore restrict our attention simply to an EMPLOYMENT/DETERMINING/EVENT/HISTORY. In order to

capture the same sort of knowledge as that represented in the state-transition network of Figure 5, we can conceive of the attribute itself as an active object, much along the lines of an abstract data type. In addition to recording a series of events of certain types (constituting the "history" of the participating object in this respect), such an abstract attribute would also behave as a monitor for events yet to occur. Since the set of possible events at any given moment is constrained by the previous event-history of the object, the analog of the state-transition network in Figure 5 would be an event-monitoring transition network. From this one can not only derive values for the secondary state-attributes of Figure 5, but one can also set triggers for automatic updates to preserve data base integrity in a natural way.

Consider, for instance, the case of two employees of a company who are married to one another. Since the SPOUSE/DETERMINING/EVENT attribute of each would be monitoring, among other things, for the death of the other spouse, the value for the derived MARITAL/STATUS attribute would be subject to automatic update for both employees. This could automatically propagate updates in tax-status, group-insurance status, etc.

3. A PARSING MECHANISM BASED UPON SI-NETWORK EXPERTISE

The usefulness of the KNET distinctions we have made so far become apparent in discussing how a conventional syntactic analyzer could guide a direct mapping from NL input into a corresponding conceptual representation, as well as for interactive construction of domain-specialized KNETs. We will consider the interpretation of the noun phrase, "NY employee". Our NL parser's lexicon, in addition to providing the syntactic processor with various grammatical/syntactic information, also contains the possible conceptual bases for the concepts (denoted by lexicon entries) which may already be present in the KNET. The lexical item NY, for instance, may be regarded as a city/state ENTITY, as a LOCATION/SITE, or as an ACTION/SITE.* The linkage between lexicon entries and the actual KNET is illustrated below for NY and EMPLOYEE:

* Note the different senses of NY in: "NY is located on the East coast," "He is a NY employee," and "No NY shipments were sent out."




FIGURE 6: LEXICON HOOKS INTO KNET

In processing an input containing the noun phrase "NY employee", the syntactic analysis routine first identifies employee as head noun and NY as potential modifier; the conceptual processor is sent this information in the form of a template which "says" that we have identified the potential concept NY/EMPLOYEE -- for which a representation does not yet exist. But certain facts are known. The syntactic analyzer will ask the conceptual unit to try and link the NY and EMPLOYEE concepts in such a way that NY can be considered an ATTR/VALUE/REFERENCE of EMPLOYEE via the schematic template illustrated in Figure 7:




FIGURE 7: CONCEPTUAL TEMPLATE

In examining the possible ways in which a path between NY and EMPLOYEE can be established to fit the template, we find that NY could be viewed as a LOCATION/SITE. We therefore try the interpretation of NY as modifying a LOCATION/SITE. Procedurally we will have to build specializations of the TO/BE/LOCATED/AT/SITE concept, as well as additional specializations of other concepts in order to create the proper connections. The pattern we are trying to instantiate is one in which the domain-object (DOBJ) has EMPLOYEE as its value/restriction. However, at the start of our matching we discover that the TO/LOCATE/AT/SITE concept has a V/R of PHYSICAL/OBJECT. The strategy at this point is to see if there is an AKO path between EMPLOYEE and PHYSICAL/OBJECT. Note the multiple vistas we have from EMPLOYEE to PHYSICAL/OBJECT (e.g., employee as a

One should note that in the structural description part of concepts, there exists not only information about how the various roles are related to one another, but also implementation-specific information on how to actually go about retrieving information from a data base in a given computer environment. The same sort of building operations which generate the concept nodes leading to the representation of NY/EMPLOYEE will tailor the information retrieval routines in the SD's -- so that for each conceptual representation of NY/EMPLOYEE we also have an implementation-bound procedure. The implementation-bound procedure is NOT the "meaning" of NY employee -- as "proceduralists" often seem to claim. Our approach makes it clear that two very different things really occur: One is the conceptual representation building process, and the other is the process which constructs an implementation-dependent representation which can be executed to check whether any NY employee exists, who they are, how many there are, etc., depending upon the intent of the original input. The latter representation, however, really is irrelevant so far as building conceptual representations are concerned. Below is illustrated the KNET after the virtual nodes have been constructed for the parsing of NY/EMPLOYEE.



FIGURE 8: INTERPRETATION OF "NY/EMPLOYEE"

Modularity has been achieved in the separation of domain independent from dependent knowledge, and the separation of syntactic, semantic (and pragmatic) expertise. No longer is the lexicon a "wastebasket" for idiosyncratic information which can't be factored out and formalized. Secondly, we have a perspicuous, comprehensible and communicable formalism which lends itself readily to conceptual/linguistic extension by a KBA who builds upon a core SI-Net for a given universe of discourse. Note that we have not gone the route of Schank and his co-workers -- who argue in favor of "semantic primitives" and the belief that all concepts can be algorithmically decomposed into their primitives (via the encoding of semantic definitions of words in terms of a small set of ACTs) for all NLS. Schank's assumptions run counter to a fairly old and well established philosophical tradition that warns against trying to arrive at a canonical decomposition of NL utterances in terms of semantic primitives. More seriously, Woods and others have pointed out that such (CD) decomposition is often undesirable (if it were possible), since necessary information which is present in the original surface encoding of the message usually gets lost. To make matters worse, the lexicons of Schank's implemented systems are enormously complex, and have never been formally characterized. The hope, then, of adapting a lexicon from one application to another seems fraught with difficulties.

Finally, we believe that the only hope for realizing truly flexible MLI's lies in the elaboration of an epistemological framework firmly grounded in the social/contractual aspects of human interaction. It should then be possible to shift to new domain applications simply through the interactive specialization of a core KNET created within this framework. We hereby dedicate all such systems to Jean Jacques ROUSSEAU.