Gouvernement du Canada
Ministère des Communications

Government of Canada
Department of Communications

Le Centre canadien de recherche sur l'informatisation du travail
Canadian Workplace Automation Research Centre

KNOWLEDGE REPRESENTATIONS: A STATE OF THE ART SURVEY

by

Kim Dalkir

Canada

KNOWLEDGE REPRESENTATIONS: A STATE OF THE ART SURVEY

by

Kim Dalkir

This report is the result of research done at the Canadian Workplace Automation Research Centre (CWARC), Laval

Department of Communications

November 1987

DD 10231991
DL 12447120

Q
387
D35
1989

ISBN 0-662-17914-5
No. Cat. Co 28-1/54-1990 E

1. SUMMARY

The major knowledge representation schemes used in the field of artificial intelligence are presented here. An overview is given of the particular strengths and weaknesses of each approach to provide a comparative evaluation of the various knowledge representations. A number of different classification schemes are also included in order to provide a general frame of reference. Samples of each representation are given in the appendix at the end.

The types of representations discussed include: semantic nets, object-value-attribute triplets, frames, scripts, conceptual dependency maps, rule or logic based methods, genetic graphs, constraint networks, concept entailment meshes and generalised concept models.

While most of the work being done in knowledge representation is within the field of artificial intelligence, particularly in knowledge engineering procedures required in the development of expert systems, other contributing fields are incorporated in this overview. Applications to a variety of domains such as psychology and education are also briefly addressed.

1. INTRODUCTION

A knowledge representation stores knowledge about a particular problem domain, much like a database does, but with several important distinctions. Knowledge processing, as opposed to date processing, involves inference and reasoning. Knowledge tends to be qualitative or at least quantifiable. This means a store of knowledge (a knowledge base) and its representation (the knowledge representation) will necessarily be imprecise, uncertain and anectodal at times. This is the nature of knowledge yet we can and do make efficient use of such knowledge every day.

Each knowledge structure may form a part of a larger system and in turn subsume one or more of its own subsystems. Knowledge is thus essentially systemic in nature. A knowledge representation must be correspondingly systemic in order to capture and represent knowledge about commonly occuring patterns in the real world. The object is to be able to capture as many generalizations about the domain as possible in order to represent increasingly larger subject areas.

A representation is a stylized version of the world. Unfortunately, a representation will also contain any referential ambiguity that is present in the real world. A good example of this is language. Words are probably the first knowledge representations we used and they are far from precise, unambiguous tools. A database, on the other hand, is necessarily unambiguous as we are forced to explicitly define all the definitions and relationships between the data objects in the structure. A knowledge structure is also more complex in that it also contains the rules and conditions that determine how and when the data will be used. This often involves insight, observation and subjective judgment.

Knowledge, unlike data, is not merely a collection of facts but a collection of meaningful information that is related by some overall structure. A knowledge representation scheme thus allows us to model this structure or framework of knowledge items required in order to make sense of a given problem area, or at least some facets of that domain. We can then retrieve the information we require, along with its consequences, within a reasonable amount of time. Knowledge can be organized and represented much like words in a language, to designate concepts and concept relationships. By making the elements and their interrelationships explicit, we render the knowledge base more accessible and therefore more comprehensible. If knowledge can be said to be a model of the real world, and a knowledge representation is a model of knowledge, then the representation is in fact a model of a model that forms a universe analogous to that of the real world.

A good knowledge representation scheme should have the following characteristics:

1. *representational adequacy* - the ability to represent all the types of knowledge and know how required for the domain,
2. *inferential adequacy* - the ability to derive new knowledge from existing knowledge,
3. *inferential efficiency* - the ability to incorporate relevant additional information into the knowledge structure when the need arises,
4. *acquisitional efficiency* - the ability to acquire new information easily.

We need to represent knowledge in order to know the relationships between concepts and to avoid inappropriate ones. This will allow us to see how the concepts are related to one another at a more global level. With such an overall view of knowledge, much like a map to gain an overall view of a city, we can make more informed choices about our or someone else's state of knowledge. This will ensure a more general understanding and, at the same time, a more coherent one, of the subject matter at hand, regardless of whether the knowledge representation is embedded within an expert system or is simply drawn out on a blackboard. By making explicit your knowledge representation, others know what you know, what you don't know and what cannot be known at present.

A knowledge representation thus formalizes the knowledge base of concepts that form a subset of the task domain or subject matter. It is a specific subset of the possible understandings that may accrue as a result. The knowledge may in fact be imperfect but this will only become evident during the course of trying to develop the knowledge representation (which is another rationale for their use in the communication process).

2. CLASSIFICATION OF REPRESENTATION METHODS

Several knowledge representation schemes have been developed in the field of artificial intelligence. Although there were representations of knowledge long before anyone tried to embue computers with a form of intelligence, the term "knowledge representation" has become almost exclusively identified with the field of AI. This overview is thus based primarily on the research of this field.

The major dichotomy in knowledge representation stems from the type of knowledge that is to be represented: static or dynamic. *Declarative* representations deal with static knowledge, mostly in the form of facts (e.g. Fido is a dog) while *procedural* representations involve processes such as how to use a particular piece of knowledge (e.g. to teach Fido to stroll a walking pace follow section 4.3 of the manual). The advantage of the declarative method is that each fact or piece of information needs to be stored only once, regardless of how many times or how many different ways in which it will be used. It is also fairly easy to add new knowledge to this type of representation. The advantage of the procedural approach is that it is much easier to represent know how, probabilistic information and heuristics (rules of thumb used to increase efficiency). Both types of knowledge are usually required for most domains therefore both types of knowledge representations will likely be required in order to capture both the knowledge and the know how involved in a given subject area.

Some common declarative representations are semantic nets, frames and scripts. Common procedural representations include rule or logic based systems such as predicate calculus. These, along with other representations, are discussed more fully in the subsequent section.

Representation schemes may also differ along other dimensions such as: deep vs. surface knowledge, qualitative vs. quantitative knowledge, approximate (uncertain) vs. exact (certain) knowledge, specific vs. generic, and descriptive vs. presriptive knowledge. Thus for every form of knowledge, a different type of knowledge representation may be best suited. In all cases, the knowledge representation chosen should be explainable and justifiable in terms of the underlying knowledge base being modelled and the representation should be restructurable and reorganizable whenever required.

Another way of categorizing knowledge representation methods would be their place on a continuum that ranges from 100 % syntactic to 100 % semantic representations (see figure 1). For example, a logic based representation would be practically purely syntactic in form whereas conceptual dependency maps would be almost purely semantic in form. The criterion used is the depth of meaning or understanding represented by the structure in question.

Figure 1. Knowledge representation spectrum

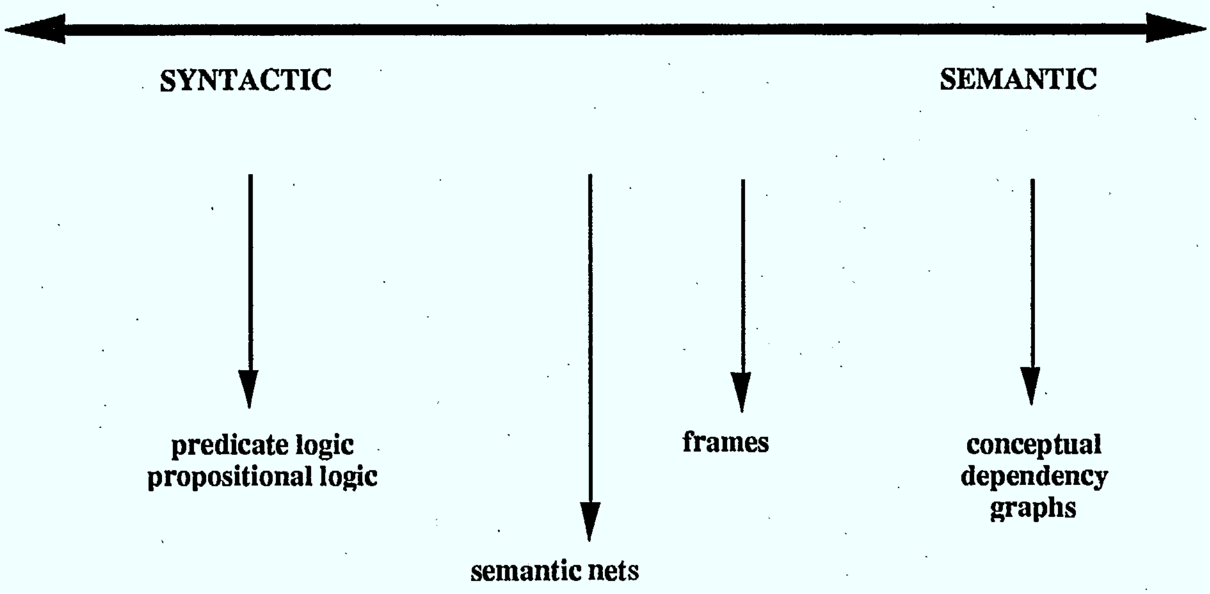

A knowledge representation, whether it be text, video or graphical form, represents the way in which the author "understood" the knowledge modelled. Thus a knowledge representation is a public representation of an individual's or a group's understanding. If the individual is a recognized expert in the field, then the knowledge modelled is expertise in that domain. We can then interact with and manipulate this representation. For example, we use words in natural languages in order to communicate. We don't directly access the concepts themselves. Knowledge representations may have great potential as a form of meta-words or tools to aid us in communicating our thoughts.

3. OVERVIEW OF REPRESENTATIONS

Models of human knowledge are often portrayed in some type of associative network. This has intuitive appeal since the human memory is believed to be organized in this manner but there are problems, particularly with the types of dependencies that can be modelled. While it is fairly simple to model conceptual relationships such as causality and conditional dependence. To convey meaning beyond connectedness, then, a graph typology that can make semantic distinctions between direct and indirect connectedness is needed. Another way of stating this is that directed graphs are needed in order to capture the complexity of knowledge; undirected graphs only give us the skeletal framework of that knowledge base.

A popular way to represent the organization of knowledge in a visual form is a structural model of some sort. Knowledge is structured as a set of facts, with each fact being related to one or more facts in a type of causal relationship. Cognitive maps, mental models maps, influence diagrams, sociograms, etc. are all similar to this approach. Having established the building blocks, i.e. the concepts, these can then be linked together to form a relational network. Linked concepts are the foundation of any knowledge base. A set of rules is usually invoked and the end result is a spatial diagram that shows all the concepts and the relationships between them.

The semantic spacing of concepts is usually represented in a variety of ways as well. Some connectors may be of the form "is a " or "has a", or more complex descriptors such as "is analogous to", "is an example of", and so on. This conceptual framework forms the basis of the domain of interest or the subject matter to be learned or taught. In general, the number of topics represented is a measure of the overall complexity of the knowledge base. For highly complex knowledge structures, one will likely need to go beyond this simple restricted network of semantic distances and use both declarative and procedural representations of the knowledge.

4.1 Semantic Nets

This is the most general form of knowledge representation. Semantic nets describe both objects and events in a declarative manner. They were originally used to represent the meaning of English words and phrases where the name was derived from. Semantic networks are also the most flexible form of knowledge representation for they imitate the nature of human memory, which is a form of associative net. (Semantic nets were originally a psychological model of human memory). Nodes in this net are connected by arcs or links that depict relationships such as "is a" or "has a". Links may also be definitional, examples, heuristics, causality and any other form of association. Both the nodes and links are usually labelled and new nodes and links can easily be defined and incorporated into the knowledge representation.

Semantic nets have the useful property of inheritance. This is the feature whereby a node automatically inherits all the properties that have been ascribed to higher more general nodes that it is attached to. The only drawback to this feature is that it becomes quite difficult to deal with any exceptions that may arise in the context of the domain.

Semantic net representations are thus very similar to indexing systems used in libraries (Dewey decimal systems). For example, node links may be bidirectional in order for intersection searches to be carried out (e.g. to find out what any given pairs of nodes have in common). The net may be nonhierarchical or nonlinear depending on the domain being modelled. This approach works particularly well for domains with well established taxonomies such as biology and natural languages.

The core of a semantic net is a set of encoded concepts and the basic unit of analysis is a conceptual graph which consists of concepts linked together in some type of association. These graphs may in turn be associated with other graphs to form the larger semantic network. Semantic nets thus form an intelligent ordering of the knowledge into an integrated descriptive model. This process, shown in figure 2 below, is in fact quite similar to the steps followed in scientific model building.

Figure 2. Semantic net methodology

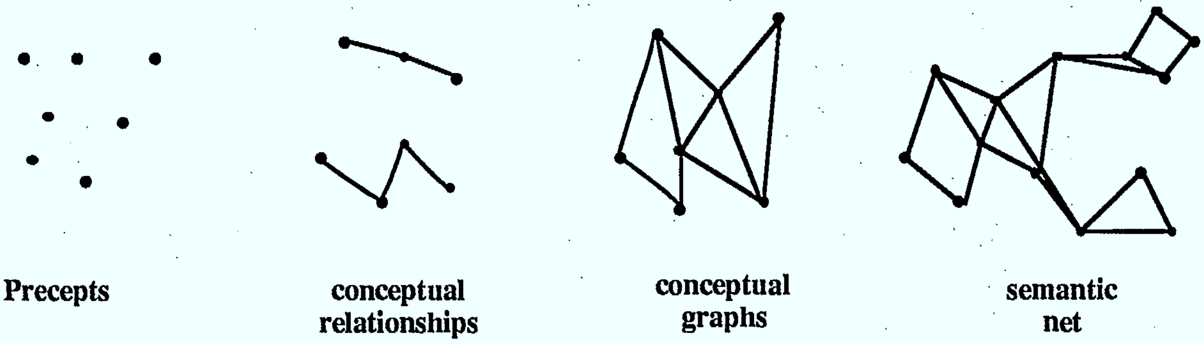

An interesting elaboration was proposed by Collins and Quillian (see Appendix) to attempt a quantitative measure of the semantic relatedness of concepts (measure of the semantic distance between concepts). They used response time to a word association exercise to assess how related two concepts were in an individual's mind. For example, when prompted with the word "bird" most people think of "wing" ("has a") before they think of "canary" ("is an example of"). One can then measure how long it takes before each word is expressed. This may in fact be an exteriorization of how we store knowledge in our memory, preferring the highest possible abstraction level and more general attributes.

In summary, then, semantic nets are useful when the knowledge to be represented is redundant, without many exceptions and that lends itself to an associative structure.

4.2 Object-Attribute-Value Triplets

In this scheme, commonly referred to as OAV triplets, objects may be physical or conceptual entities, attributes may be any general characteristics or properties of these objects, and values may be any specific characteristic of the object. The value of an attribute specifies the exact nature of a particular object or situation. The process of giving a general attribute a specific value is called "instantiation". For example, the object apple has the attribute color and a Macintosh apple has the value red for that attribute.

An OAV representation is really a specialized form of the semantic net where only "is a" and "has a" relationships are permitted between objects in the representation and there are only three types of nodes (object, value and attribute).

Both static and dynamic knowledge can be represented using this scheme. Static knowledge is the case where unchanging objects have general attributes. Dynamic instances of that object, however, can and do vary from case to case. The second important feature of OAV representations is the manner in which objects can be ordered and related to one another. The graphs may be treelike in form, with the topmost object as the "root" which is used as the starting point. These trees may also have a dynamic nature of their own. Objects other than the root object can also have multiple instances. In fact, the object trees often become entangled so that subordinate topics end up being related to more than one higher-level object so that they can inherit properties from more than one source.

OAV representations can be used to handle uncertain or incomplete knowledge using certainty factors. The latter refers to quantifications of the degree of certainty we can attribute to each fact or inferential conclusion reached based on those facts. The best known example of this can be found in the expert system MYCIN, which provides diagnosis of bacterial infections along with a certainty factor (e.g. patient A has infection XYZ with a certainty factor of 0.70 or 70 %). These are not probabilities in the statistical sense (although they can be) but rather measures of subjective belief or subjective preferences (e.g. I like apples; green apples 0.3 and red apples 0.7).

Finally, some knowledge systems are built around a single group of facts or objects. In these cases, object-attribute pairs are sufficient to capture all the knowledge required. There are no values required for attributes since there is no change and therefore, no inheritance.

4.3 Frames

The frame approach can describe objects in terms of a collection of attributes possessed by them. A declarative form of knowledge representation is used in this scheme. This is a general method that works well for the representation of complex objects that may be viewed from a number of different perspectives.

There is some evidence that humans do not process information from scratch but instead build on existing knowledge structures. Thus, when confronted with a new fact, we do not build up a new knowledge structure to describe this situation but instead think back to similar situations in order to decide how best to deal with the present, new situation. We have available in our memory, a large collection of experiences which can be applied to new knowledge. A single frame is usually one such representation of this type of common knowledge. A frame typicallty describes a class of objects using a collection of slots that describe the various aspects of objects. These slots may contain conditions with assigned default values or they may have procedural information on which frame to go to next, for example.

Selected frames may be grouped together to form a frame system in order to represent the different points of view on the same object. For example, one can study a subject from a variety of different perspectives, depending on our own particular backgrounds. This allows us to infer unobserved facts about new situations and to represent typical instances of the concepts they represent. Partial evidence can be used in this manner to initially select a frame. This frame can then be instantiated as more information concerning the specific instance or current situation is obtained. If no appropriate values can be found, then a new frame is selected, based on the degree of similarity to the original frame. This strategy ensures a common thread of continuity is maintained (refer to similarity nets in the Appendix).

Frames are thus data structures used to represent a stereotypical situation. Several kinds of information are attached to each frame, such as how to make use of that frame (procedural knowledge) and what the other related frames are (declarative knowledge). The organization is very similar to that of a semantic network - in fact, frames and semantic nets are both "frame-based systems" in more general terminology. These types of knowledge representation schemes are particularly useful for domains where expectations play a strong role, since they can easily be expressed using frame relations.

Frames allow for a richer representation of knowledge but they are also more difficult to develop because of this complexity. They have a great deal of flexibility in that dual semantics (declarative and procedural knowledge) can be accomodated. The procedural items relate to instructions for determining or computing an entry. Information from other slots or other frames may also be incorporated here. Thus both forms of knowledge can be represented within a single frame which facilitates any transformations between the two. This also increases the power, generality and popularity of the frame knowledge representation.

4.5 Scripts

This is a more specialized form of knowledge representation which is used to represent common event sequences. Scripts describe a stereotypical sequence of events within a particular context. These are useful because in the real world, there are often patterns in the occurrences of events. The initial event represents the entry condition for that sequence to be invoked. For example, by deciding to eat out (initial event), a whole sequence of fairly standard events will ensue: deciding where to eat, how to get there, getting there, ordering a meal, eating a meal, paying for the meal and returning home. Scripts thus arise from a string of causal relationships between events.

Scripts are useful in modelling the occurrence of events that are not explicitly stated, in order to limit redundancy. In our example of eating out, one doesn't usually state that the meal was paid for: this is inferred because we are familiar with the context of the restaurant script. In a script, it is not necessary to state each fact explicitly once the sequence has been formalized. Scripts are also a useful way of building a single coherent interpretation from a collection of observations. They can be used to focus attention on only the more unusual events that arise (similar to exception reporting, i.e. you only want to be told when something is *not* going according to plan).

4.6 Conceptual Dependency Maps

This is another specialized form of knowledge representation which is used to represent relationships between components of an action. For example, conceptual dependency can be used to represent the meaning of language in such a way as to facilitate inferences to be made. It can render concepts independent of the original language used. The set of allowable dependencies between the conceptualizations expressed are then identified for these primitives. The dependency maps are themselves conceptualizations and can thus be used as components of larger maps. Beliefs may also be accomodated in this representation scheme.

Conceptual dependency maps facilitate reasoning due to the following properties:

1. fewer inference rules required and these need be expressed only once in the representation,

2. many of the inferences are already contained in the representation itself (implicit rules), and

3. the initial structure has gaps which serve to focus attention.

The only drawback to this representation methodology is the need to identify the set of primitives. This may be difficult or impossible to do for some subject areas or it may not necessarily be desirable as it may render the structure too redundant and inefficient. On the other hand, no special expertise or training is required to create conceptual dependency maps. The basic entity is the conceptualization of some event to be represented and assert which events happened. Thus conceptual dependency maps are best suited to representing activities of some sort. They are also more amenable to representing more abstract forms of knowledge.

4.7 Rule or Logic Based Methods

Rule-based knowledge representations depict relationships, usually in the form of some type of IF-THEN rules. There may be certainty factors attached to these rules in order to deal with uncertain knowledge. Rules may also be variable in that they are matched up with the facts in a sort of lookup up (decision or match table). This allows inference chains to be invoked and the "best" rule for each situation to be selected (refer to figure 3 below).

Figure 3 Pattern matching

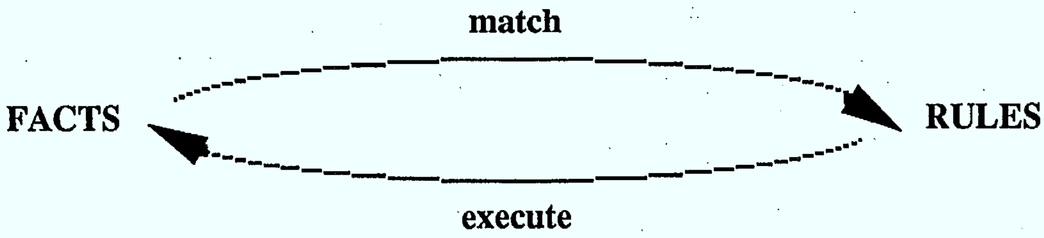

Most rule-based knowledge structures operate by choosing, on each cycle, the first rule that matches. This requires an implicit ordering of rules in the database since some rules are more likely to apply than others (i.e. more likely to lead to a good solution). There may be intermediate or partial solutions generated during the course of this process. The associated belief or disbelief index of each such assertion can also be recorded and revised using the same certainty factors.

Rules are thus a formal way of representing knowledge and recommending directions or strategies to follow in using this knowledge. This form of knowledge representation is most appropriate when the domain knowledge results from empirical associations developed through extensive testing and/or extensive experience.

Logic-based knowledge representations are powerful, rule-based schemes that allow new knowledge to be derived from old, using mathematical derivations. The two major categories are propositional logic and its extension, predicate calculus. Other logic systems allow various degrees of certainty, inference, heuristics and beliefs.

Propositional logic deals with true-false values whereas predicate logic allows the calculation of the actual degree of "truth" in a given proposition. A predicate is simply something that asserts a fact about one or more entities. The atomic formula for a generic rule is:

IF (premise)...THEN (conclusion)

where the premise and conclusion may be multiple conditions. The atomic formula for a predicate logic representation is in the form of:

(predicate -- terms) ... assertions.

For example, instead of saying, "the ball is red", in predicate logic it is stated in the form (ball, red). The advantage of the latter approach is that it can support the inference techniques of chaining or resolution. In other words, it lends itself more to domains where a great deal of knowledge must be derived instead of being explicitly stated. However, note that the logic guarantees only that conclusions are true if the premises were true in the first place ("garbage in garbage out" still applies).

The other forms of logic representations include nonmonotonic reasoning, fuzzy logic and probabilistic reasoning. Nonmonotonic logic differs from the ordinary run-of-the-mill logic in that statements can not only be added but they can also be deleted from the knowledge base. This is required for most belief systems since beliefs are constantly being revised. For example, it is asserted that the best way to get to the airport is by car until you listen to the weather report whereupon you decide the best way is by bus. This type of reasoning can be accomodated in nonmonotonic reasoning much more easily than by propositional logic. In *monotonic* logic, the statements that are known to be true can only increase over time. *Nonmonotonic* or default reasoning allows the number of true statements to both increase and decrease. The latter does entail a problem, however, in that consistency checks must be carried out each time beliefs are revised.

Probabilistic reasoning is used to represent likely but uncertain knowledge. It is used in three basic types of situations: when the relevant world or subject matter is really random in nature, when there is not enough date, and when we do not fully understand the relevant domain. A variety of statistical inference methods, such as expected value theory and Bayes' theorem may be used in depicting the relationships between uncertain knowledge items in these forms of representations.

Knowledge may also be represented using fuzzy sets. These consists of heuristics or rules of thumb, judgmental rules and plausible reasoning as well as facts that have varying truth values. Knowledge diagrams are in fact models of fuzzy sets by their very nature. There can be a variety of functions between different fuzzy sets or universes of discourse. Rules of inference are therefore needed to construct fuzzy set categories which are typically of the form: "sort of tall", "moderately easy" and so on. A knowledge representation can display fuzzy sets and the relationships between them. These schemes are appropriate for knowledge that varies as a continuous variable and for objects that have properties that are not easily quantifiable.

Logic representations are best shown using Venn diagrams, where overlaps designate sets of objects with common characteristics. In propositional logic, the propositions are linked together by connectives such as "and", "or", etc. to form compound propositions. This is a useful way of representing rules for the propagation of truth in statements and also in sets of inference rules. Predicate logic has the object as the elementary unit in the representation. Predicates are statements of the object which may be true or false and they may address more than one object. These predicates are also linked together by connectives such as "and", "or", and so on. Logic knowledge structures are thus best suited to relationships of this type. The artificial intelligence language PROLOG is perhaps the best known example of the predicate logic approach to knowledge representation schemes. Database query applications have been the most successful PROLOG implementations to date.

4.8 Genetic graphs

The genetic graph is a particular framework that was developed for student modelling using analogies. Analogies are important in any knowledge representation scheme because they help to organize knowledge. We tend to build on existing knowledge whenever possible which means we make extensive use of analogies in learning and communicating our thoughts. Genetic graphs capitalize on the use of analogies to derive student models for intelligent tutorial systems.

A good student model should be able to assess the following features:

1. knowledge (facts, concepts, procedural skills) that the student has mastered,
2. items of knowledge that the student has not mastered,
3. knowledge we believe was mastered but for which we lack any evidence of mastery,
4. misconceptions or procedural deviations, and
5. any learning style preferences.

The advantage of the generic graph is that it is very flexible, it has been used successively in many domains and it makes use of analogy links between knowledge items or concepts. A genetic graph is essentially a directed graph with nodes to represent knowledge (facts, knowledge, skills, deviations, etc.) and nodes to relate these in a variety of ways (generalization, component, analogy, deviation, corrections, etc.). The graph may be a multidimensional one, with various links to indicate the relative difficulty of each concept and its subskills. In addition, nodes along with their links may be formed into islands to represent a single skill or body of knowledge concerning a particular concept. Post and pre links are another elaboration, that may be used to depict prerequisite relationships.

A sample genetic graph is shown in the appendix. Note that there can be analogies both between nodes and between subgraphs. There can also be any type of knowledge representation scheme at each of the nodes (e.g. frames, scripts, etc.). The goal is to have a dynamic model of the student in order to be able to accomodate as many different student users as possible. The ability to incorporate any other form of knowledge representation serves to increase the flexibility of this approach.

4.9 Constraint Networks

This is a straightforward representation of objects and their relationships where objects are represented as nodes while their relationships are expressed in the form of constraints. Two objects are related if their nodes are connected through some set of constraints. Node descriptors are special kinds of relationships used to represent parts, attributes and properties of objects. Thus constraint networks provide a single mechanism for representing all three types of knowledge representations (whereas other schemes would require a distinct structure for each).

A node descriptor may contain either an object or a collection of objects. Object collections allow intelligent cooperation between data structures that hold multiple objects of the same type, or real world objects which often have multiple-valued attributes. This organization can then be indexed and ordered so that the entire collection can be accessed and implemented as one data structure.

Constraint networks express a program but specify it in a non-procedural form. That is, no algorithms are required in these networks. An example is shown below in figure 4, for a square root constraint network. In this case, it is only stated that B time B is A and that both A and B must be greater than zero. Processing always begins at the nodes with computations to be performed in order to find these values first. The constraint network, like the semantic net, allows property inheritance and the additional ability to copy an entire network structure.

Figure 4. Square root conxtraint network

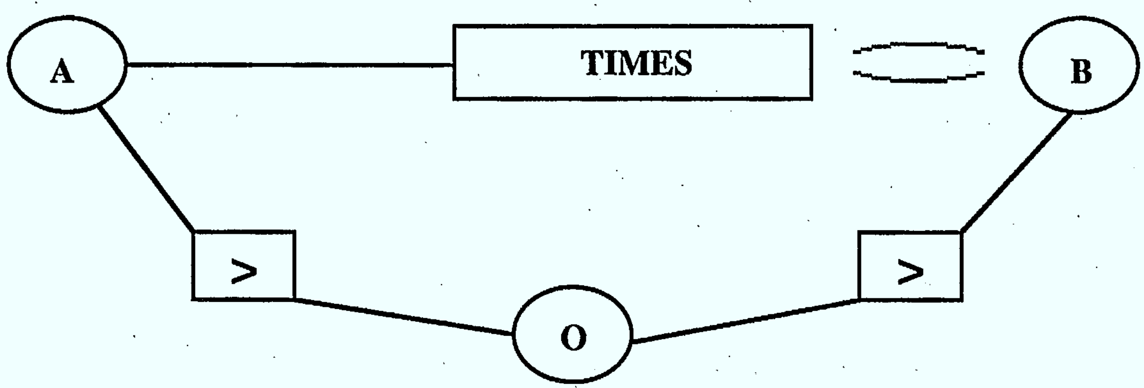

A constraint expresses a computational relationship among the nodes it connects. A complex constraint is thus one that expresses the relationship in a separate network of constraints. In this way, a single constraint can be used to embody an entire network of knowledge and increase its abstraction level (meta knowledge representations).

Constraint structures can also be transformed, modified and inferred through the use of pattern matching techniques, similar to those of rule or logic based methods. This representation scheme is the best way of representing nonprocedural specifications that must be transformed into procedural ones.

4.10 Concept Entailment Network

This form of knowledge representation was developed within the context of Conversation Theory and the protolanguage Lp, by Gordon Pask. A conceptual entailment network is useful in domains where in order to understand one topic, one must also understand other related topics. This is consistent with the systemic view of knowledge and knowledge structures. Relationships are essential features which serve to characterize knowledge systems. Thus knowledge may be depicted as a relational network or entailment network. This network identifies not only the topics and their relationships but also the cognitive processes, structures and clusters of procedures associated with them. A visual representation of the domain is provided which enables one to "see" the scope of the subject.

This type of network is more lattice-like in structure, as opposed to the more conventional hierarchical or tree-like knowledge representation schemes. One can think of a mesh as a type of net, with knots representing the nodes. Any one knot or node can be temporarily held up above the others to denote a target concept (to be accessed or learned). Topics can thus be studied at a variety of levels and from a number of different perspectives. The overall structure is an integrated, higher order form of knowledge. If subjects can be made aware of the nature of the form of their knowledge about a particular topic in the form of a mesh, then one could use meta processes on "knowing about knowing" to conduct a lesson, a conversation or to drive a search strategy. Each mesh thus represents one way a particular individual thinks about a certain domain, including any biases, misconceptions, etc. he may have. These personal knowledge structures become publicly accessible, once they are in the form of an explicit mesh.

A concept in this form of knowledge representation is a mental or cognitive representation of an object, idea or feeling, a mental process through which we create and recreate mental images of these objects. Concepts have a contextual nature in that they tend to form and exist in clusters. A topic, then, is a computer representation of a concept and a concept map is a graphical representation of the relationship between concepts showing semantic closeness. The latter refers to the perceived ordering of concepts with regard to other, related concepts in order to learn the target concept (e.g. temporary prerequisite hierarchies). This is in fact an instructional design superimposed on the knowledge representation.

Entailment meshes are thus well suited for flexible, dynamic representations of a subject matter that isto be taught, learned or somehow communicated to others. Meshes are easily modified, decomposed and joined together to form higher order structures. Finally, each node can accomodate a variety of declarative and procedural knowledge representations, which further increased the utility of the mesh.

4.11 Generalized Concept Representation

The Generalized Concept Model uses a hybrid frame and logic based knowledge representation for domain objects. The advantages of frames are the naturalness of organizing knowledge as taxonomies of generic frames which are then filled up by particular instances. The main disadvantage is a lack of formal semantics. By combining frames with logic, the Generalized Concept model makes up for this deficiency. Objects in the domain are logically associated with their descriptors and there is successive refinement of the knowledge representation using frame-like conceptual descriptions. Thus this scheme provides for a type of controlled property inheritance in the network (no longer autonomous as with semantic nets).

There are two types of conceptual objects possible: concepts and relationships. A concept is a collection of entities within a context and with similar content. The information content of a concept consists of a finite number of attributes including mappings from or relations between the concept to other domains. Each attribute is associated with a single domain. Relationships are links between concepts or sets of concepts and may be of the "is a" or "has a" form. In addition, the links themselves may have attributes. Generalized concept representation schemes thus appear to combine a number of useful features found in other approaches.

BIBLIOGRAPHY

Chandrsekaran, B. and Mittal, S. (1982). Deep vs. compiled knowledge approaches to diagnostic problem solving. Paper presented at the AAAI conference, 1982, Pittsburg, Pennsylvania.

Collins, A. and Quillian, M. (1969). Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behaviour, 8, 240-247.

D'Ambrosio, B. (1985). Expert systems - myth or reality. Byte, January, 1985, 275-282.

LaGiadice, D. and Scaruffi, P. (1986). Knowledge modules vs. knowledge bases : a structure Xfor representing the granularity of real world knowledge. Proceedings of the 6th Canadian conference on AI, May 1986, pp. 220-224.

McAleese, R. (1985). Some problems of knowledge representation in an authoring environment: exteriorization, anomalous state metacognition and self confrontation. PLET, 22, 4, 299-306.

McCalla, G. and Cercone, N. (1983). Approaches to knowledge representation. IEEE Computer, October, 1983, pp.12-18.

Michaelson, R., Michie, D. and Boulanger, A. (1985). The technology of expert systems. XByte, April, 1985.

Moody, H. (1986). AI: current reality, future promise. Info. Strategy, Winter, 1986.

Pask, G. (1976). Conversation Theory. Holland : Elsevier.

Pearl, J. and Paz, A. (1986). On the logic of representing dependencies by graphs. Proceedings of the 6th Canadian conference on AI, May 1986, pp.94-98.

Pracht, W. (1986). GISMO : a visual problem-structuring and knowledge organization tool. XIEEE SMC,16, 2, 265-270.

Stefik, M., Aikins, T., Benoit, J., Birnbaum, L., Hayes-Roth, B. & Sacerdoti, E. (1982) The Xorganization of expert systems: a tutorial. AI, 18, 135-173

Thompson, B. and Thompson, W. (1985). Inside an expert system Byte, April 1985, 315-X330.

Wellsch, K. and Jones, M. (1986). Detecting analogous learning. Proceedings of the 6th Canadian conference on AI, May 1986, pp. 17-23.

Zadeh, L. (1983). Common sense knowledge representation based on fuzzy logic. IEEE XComputer, October, 1983, pp. 61-65.

APPENDIX : DIAGRAMS OF DIFFERENT FORMS OF KNOWLEDGE REPRESENTATIONS

a) semantic Nets

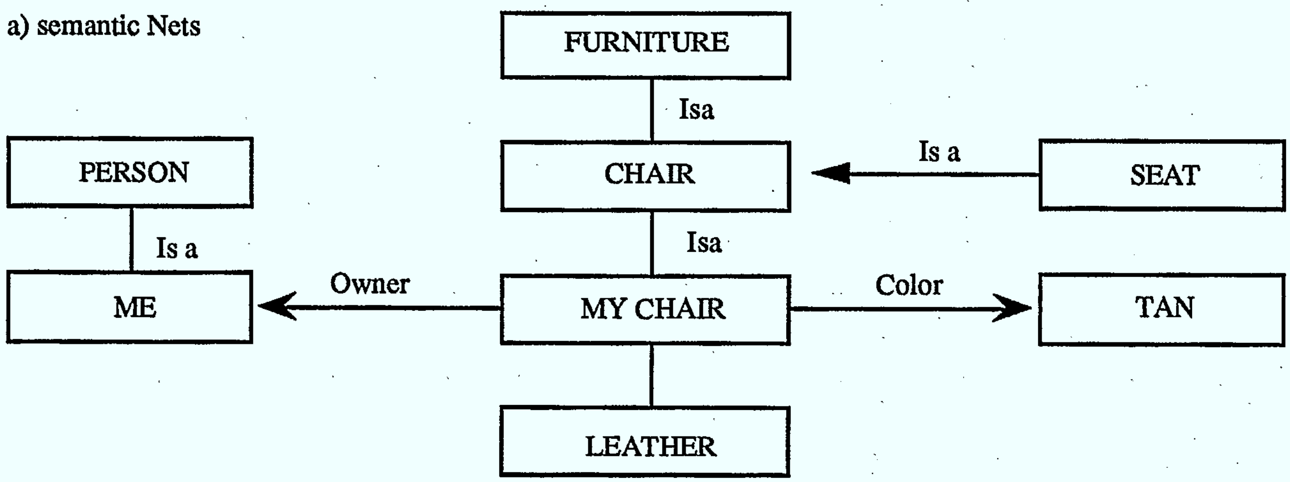

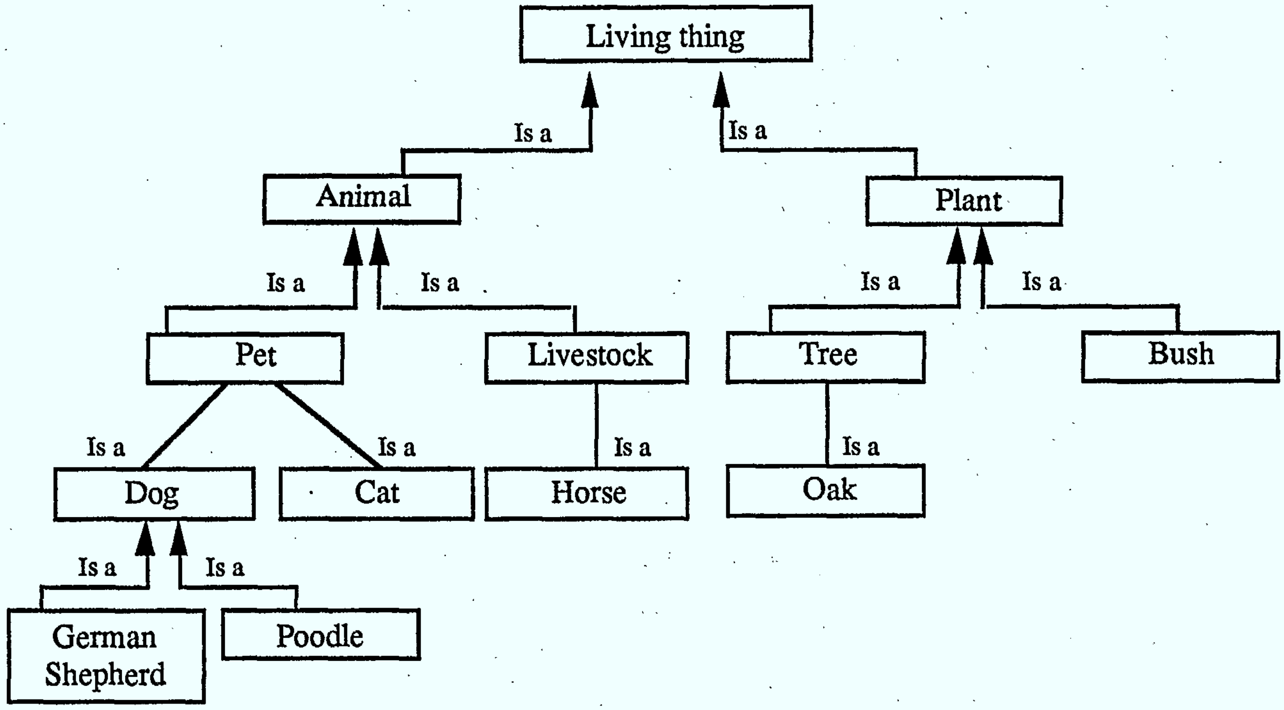
Living thing
Is a
Is a
Animal
Plant
Is a
Is a
Is a
Is a
Pet
Livestock
Tree
Bush
Is a
Is a
Is a
Is a
Dog
Cat
Horse
Oak
Is a
Is a
German Shepherd
Poodle

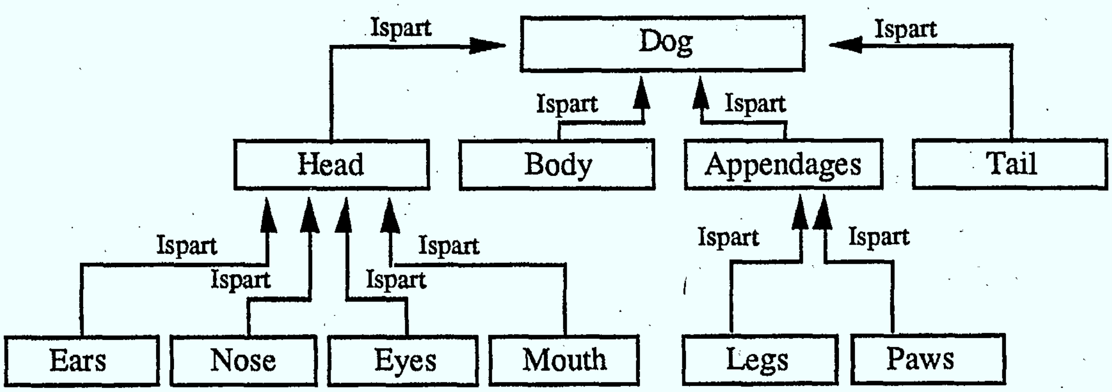
Ispart
Dog
Ispart
Ispart
Ispart
Head
Body
Appendages
Tail
Ispart
Ispart
Ispart
Ispart
Ispart
Ispart
Ears
Nose
Eyes
Mouth
Legs
Paws

b) Object-Attribute-Value Triplets

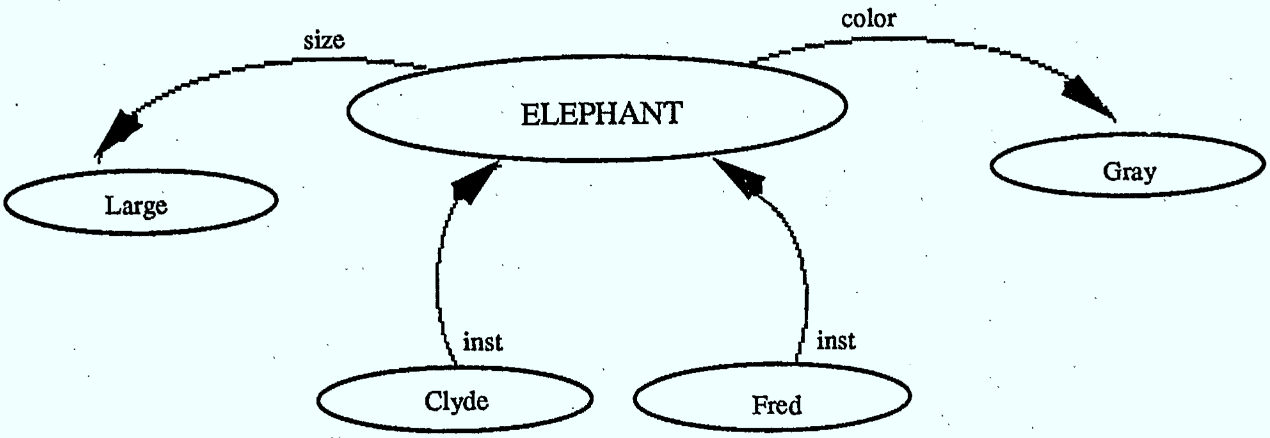

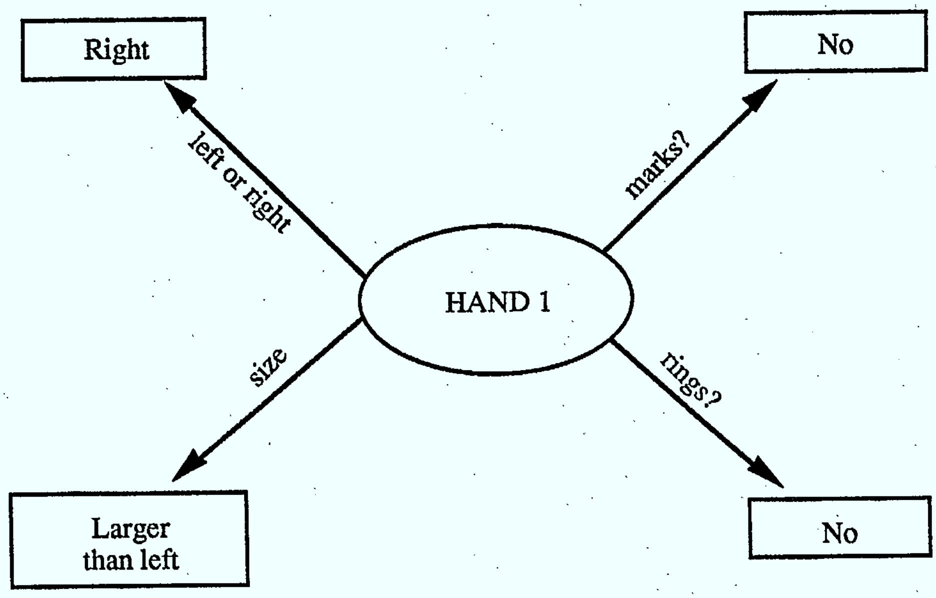

c) Frames

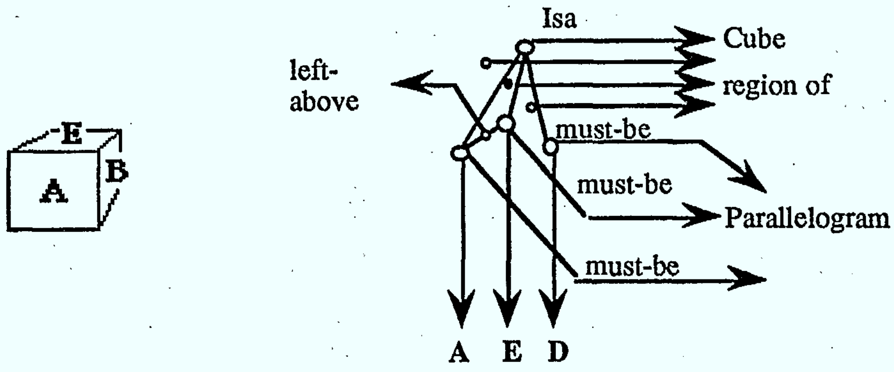

A Frame Representing one view of a cube

d) Conceptual Dependency Graph

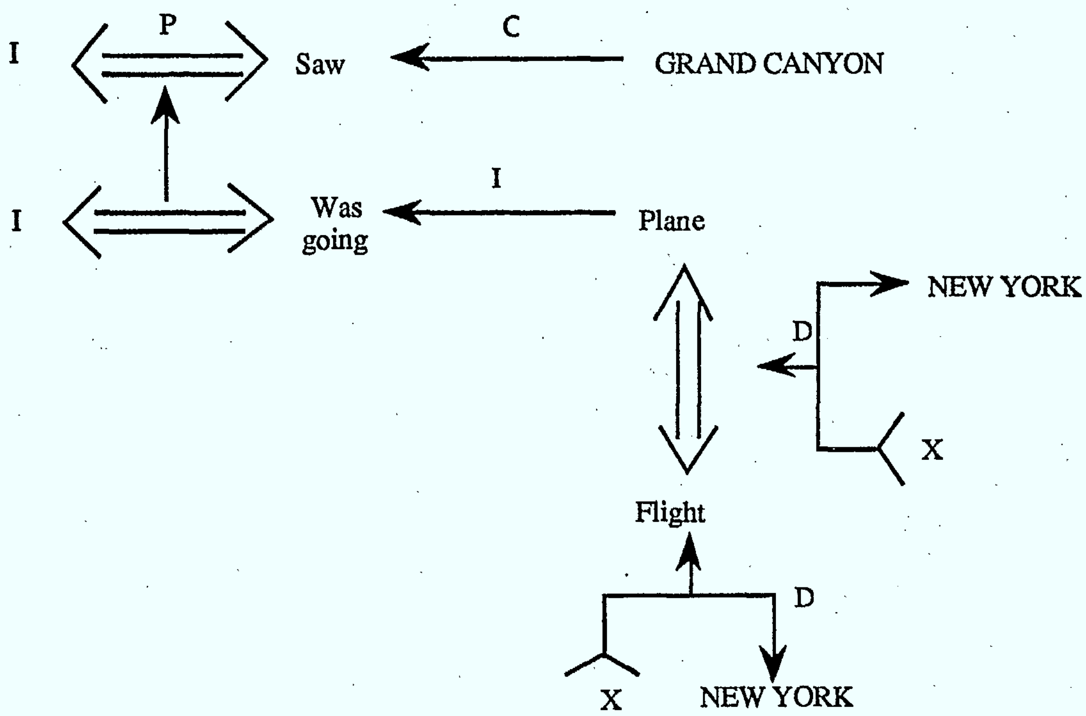

A conceptual dependency graph for the phrase.
"I saw the Grand Canyon from the airplane on my way to Chicago"
(Schank, 1971, reproduced with permission).

e) Genetic Graphs

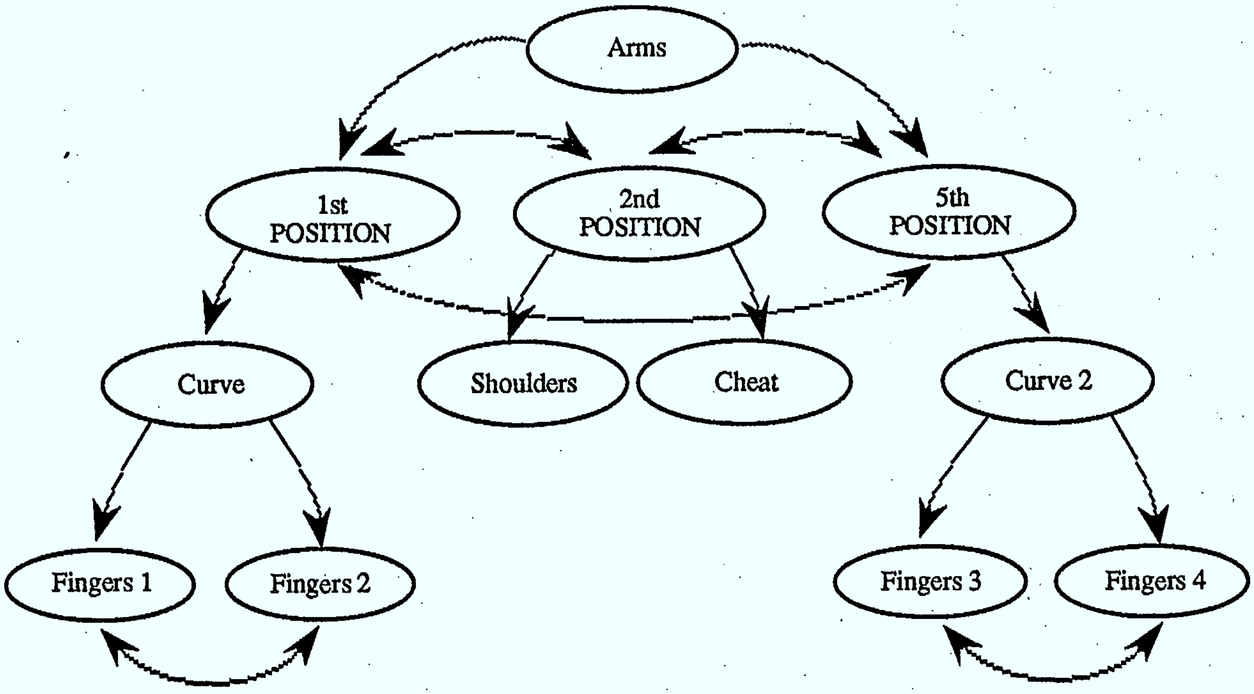

f) Similarity Nets

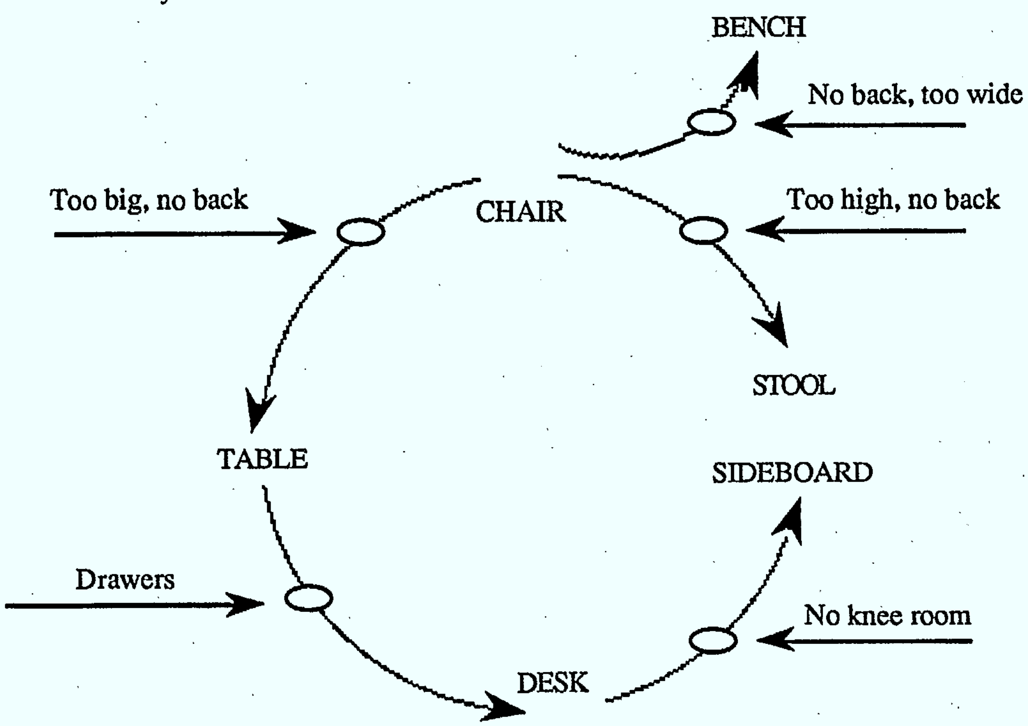

Pour plus de détails,
veuillez communiquer avec :

Le Centre canadien de recherche
sur l'informatisation du travail
1575, boulevard Chomedey
Laval (Québec)
H7V 2X2
(514) 682-3400

For more information,
please contact:

Canadian Workplace
Automation Research Centre
1575 Chomedey Blvd.
Laval, Quebec
H7V 2X2
(514) 682-3400