

# IRIS: A SYSTEM FOR THE PROPAGATION OF INFERENCES IN A SEMANTIC NET

Michael Trigoboff  
Department of Computer Science  
Rutgers University  
New Brunswick, N. J. 08903

Casimir A. Kulikowski  
Department of Computer Science  
Rutgers University  
New Brunswick, N. J. 08903

## Abstract

IRIS is a system which uses a semantic net and link-associated decision tables to reason about medical problems in a predominantly bottom up, event driven manner. The nodes in the semantic net represent medical states (e.g. hypertension, diabetes) which may be true of the patient, while the links represent relations between the nodes (e.g. causes, predisposes, characterizes). The truth of any node with respect to the patient is represented by assertions associated with that node. The decision tables control a process called propagation in which an assertion about a given node can result in assertions about neighboring nodes. IRIS uses the propagation process to perform all necessary inferences.

## 1. Introduction

Several artificial intelligence approaches to medical consultation have been developed in recent years (Pauker, 1976) (Pople, 1975) (Rubin, 1975) (Shortliffe, 1974). A method for representing diseases and their treatments in the form of a causal-associational network, and a consultation program that reasons on the basis of such a model has been developed within the Rutgers Research Resource on Computers in Biomedicine (Kulikowski, 1976) (Kulikowski, 1971) (Weiss, 1974).

In order to explore new representations for medical knowledge on an experimental basis, we have recently developed a new system called IRIS. IRIS has a uniform representation for medical knowledge which permits specification at any level of abstraction. This representation is not tied to any specific conceptual relationship (e.g. causality), but is capable of representing a variety of structures and relationships. A uniform data structure for medical knowledge also facilitates the construction of mechanisms for explaining IRIS's behavior with respect to the medical knowledge it is based on.

Medical consultation systems need the ability to represent two distinct kinds of medical knowledge which can be referred to

as general medical knowledge, and specific medical knowledge. General medical knowledge includes knowledge of diseases, physiological states, treatments, and their interrelationships. This is the knowledge on which a system bases its decisions. Specific medical knowledge is the knowledge the system has about the patient currently under consideration.

Medical reasoning uses general medical knowledge as the basis for drawing conclusions from the facts about an individual case. In designing IRIS we wanted to create a system which could easily model a large variety of patterns of medical reasoning. We achieved this goal by implementing a system in which assertions called ISPECs (information specifications) could propagate through a semantic net under control of link-associated decision tables.

## 2. Overview of IRIS

### 2.1 The Semantic Net

IRIS represents general medical knowledge using a semantic net. A semantic net (Woods, 1975) is a directed graph in which the nodes represent concepts, and the links represent relations between them. Conventionally, any two nodes connected by a link are connected in the reverse direction by the link's relational inverse. In IRIS's semantic net, the nodes represent clinical concepts which can in some sense be said to be "true" of the patient, or of the clinical situation (e.g. diseases, symptoms, treatments, etc.). The links represent various relations between these concepts (e.g. "causes", "treatment for", "associated with", etc.).

### 2.2 ISPECs

Specific knowledge is represented by data structures called information specifications, or ISPECs. Each node in the semantic net can have a list of associated ISPECs. These ISPECs specify in what way the concept represented by the node is true of the patient. There are slots in ISPECs for strength of belief, a time notation, a side notation (RIGHT, LEFT, or NIL), and modifiers which can serve to further specify the concept represented by the node. At any point, the set of existing ISPECs represents IRIS's total knowledge about the patient.

### 2.3 Decision Tables

An IRIS decision table can be thought of as a set of rules which determines propagation of ISPECs across the link it is associated with. Each rule specifies a way in which the link's source node (the node on the "near" side of the link) must

be true of the patient. If a rule is satisfied, (by some set of ISPECS associated with the source node), it then specifies the creation of an ISPEC for the target node (the node on the "far" side of the link). The result of propagation across a link is a set of ISPECS for the target node, each one produced by a satisfied rule.

#### 2.4 Propagation

When an ISPEC is generated for a node, all the links emanating from that node are potential candidates for propagating the effects of that ISPEC to neighboring nodes. Only those links which have an associated decision table can actually support propagation. For each such link, the decision table is evaluated, and the ISPECS produced are added to the ISPECS list for the node at the far end of the link. The process then repeats recursively, potentially producing a wave of propagation across the semantic net. Each new ISPEC is like a droplet which, when added to our pool of knowledge, causes a ripple which can propagate any distance through the pool.

#### 2.5 Cones

Consider a semantic net node representing a disease, and a set of nodes representing the symptoms of that disease. Whenever an ISPEC is produced for one of the symptom nodes, the truth-status of the disease node may be affected. Furthermore, these effects are generally not independent. What we need in this situation is a way for the symptom nodes to interact in propagation to the disease node.

In IRIS, this is accomplished by associating the same decision table with all the links from symptom nodes to the disease node. We call such a configuration a propagation cone. It is usually visualized with the disease node at the apex, and the other nodes occupying the cone's base.

Cones can also be created in which the direction of propagation is from the apex to the base. In this case, a single decision table is associated with all the apex to base links. We then have a structure which can be characterized as a "scatter cone".

A simple semantic net link connects two nodes. Propagation is defined by the decision table associated with that link. A cone is a generalization of this basic structure. With cones, there can be a set of nodes rather than a single node on one side of the structure.

#### 2.6 QUESTIONS

All of IRIS's specific medical knowledge must ultimately be based on information provided by the user. There is nothing we can conclude about a specific patient a priori. Accordingly, we need a mechanism by which specific medical knowledge can be entered. QUESTIONS provide such a facility.

An IRIS QUESTION consists of the following: a set of strings to output to the user, a set of checking functions to ensure that the user's input is in the proper form, and a set of mappings from permissible user inputs to the various actions IRIS can take. These actions include setting system variables, specifying questions to be asked next, and generating ISPECS. The set of ISPECS produced in this way comprises that part of IRIS's specific medical knowledge obtained directly from the user.

#### 2.7 Vines

Once a consultation system has come to a set of conclusions about a patient, it must be able to communicate them to the user. IRIS does this by generating an english language string for each semantic net node which represents a conclusion. This string expresses how the concept represented by the node is true of the patient. We refer to these strings as "vines" because they can be visualized as festooning the semantic net.

There are two ways to generate a vine for a node. The default (taken when the node carries no explicit specification for vine generation) is to create a vine for the node using a standard procedure. This procedure produces a string of the form: {modifiers, time, node-name, strength of belief}. An example would be: "severe sudden past increased intraocular pressure [SB = .85]". Alternatively, we can associate a decision table with a node. This decision table can then specify different ways of generating the vine, or different pre-stored vines. The choice among the various possibilities can be based on the truth-status of any node or set of nodes.

#### 2.8 Explanation

A clinical consultation system must be able to explain or justify its decisions in terms acceptable to experts in its application domain. Systems that lack this ability have not gained wide acceptance even when they attained a satisfactory level of correctness in their decisions. Such an explanation should refer to the general medical knowledge on which the decision was based.

In IRIS, this is easily done. Suppose that there are two semantic net nodes, "injury" and "pain", which are connected by a "causes" link (i.e. "injury causes pain") which has an associated decision table. If IRIS is told that the patient has an injury, it can then conclude (via propagation) that the patient is in pain. When asked to explain this reasoning, IRIS can reply: "I conclude that pain is present because I know that there is an injury, and I know that injury causes pain."

### 3. ISPECs

As stated above, an ISPEC is an assertion that the concept represented by a semantic net node is true of a patient, or of the clinical situation. ISPECs are IRIS's only representation for specific medical knowledge. There are slots in each ISPEC for specifying details about how the concept represented by a node is true. We will now describe these slots, and the currently implemented data structures that fill them.

#### 3.1 Node

The NODE slot of an ISPEC specifies which semantic net node it is an ISPEC for. This slot is usually not necessary, but is needed for propagation in scatter cones (see section 2.5), and for context dependent decision tables (see section 4.2).

#### 3.2 Strength of Belief

The strength of belief (SB) slot of an ISPEC specifies how much IRIS believes the assertion represented by that ISPEC. The possible range is from full belief to full disbelief. IRIS has been designed so that the mechanisms for manipulating numerical strength of belief measures are easily changed or replaced. Currently, IRIS uses the strength of belief representation developed by Shortliffe for MYCIN (Shortliffe, 1974). This is described in greater detail in (Trigoboff, 1977).

This representation for strength of belief consists of two numbers: a measure of belief (MB), and a measure of disbelief (MD). Each of these can range from 0 to 1 with 0 representing no (dis)belief, and 1 representing full (dis)belief. The MD is subtracted from the MB to obtain the certainty factor (CF). The CF ranges from -1 (full disbelief) through 0 (no opinion) to 1 (full belief). The SB slot of each ISPEC is filled with a two-element list of the form (MB MD). Thus, IRIS's strength of belief in an ISPEC is derived by using the contents of its SB slot to generate a CF.

#### 3.3 Side

The human body is bilaterally symmetrical, and contains many paired structures. An assertion about any such structure must therefore specify which one it refers to. Since IRIS is to perform as an ophthalmological consultant, it needs the ability to deal with this issue.

There are basically two ways to approach this problem. Either we can have semantic net nodes that represent "pain in the left eye" and "pain in the right eye", or we can have one node representing "pain in the eye" and rely on some notation in the ISPECs to tell us whether it is the left or the right eye. In IRIS, we have chosen the second course, avoiding an otherwise unnecessary duplication of semantic net nodes. Each ISPEC has a SIDE slot which can contain either LEFT, RIGHT, or NIL.

Not every node in the semantic net refers to a duplicated structure. "Diabetes" is a concept which is not side-related. Nodes of this sort are marked as SINGLE nodes. Any ISPEC for a SINGLE node must have a SIDE slot value of NIL. Conversely, any ISPEC for a non-SINGLE node must have a SIDE slot value of either LEFT or RIGHT.

#### 3.4 Time

Clinical situations evolve through time, and the configuration of assertions describing a patient changes as new events become known to the consultation system. Any assertion must therefore be associated with a specification of the time interval it applies to. In IRIS, each ISPEC has a TIME slot. TIME slots are filled with a list of two dates of the form (from to). The first is the date the ISPEC initially became true of the patient, and the second is the final date the ISPEC was true.

The creation of ISPECs is controlled by means of create-expressions which specify characteristics for the ISPEC to be created. Values for the TIME slot are generally specified in create-expressions as one of the following: past, present, past-or-present, future. When an ISPEC is created from a create-expression, the current date (obtained from the system) is used to translate these specifications into IRIS's internal time representation (the list of two dates).

The time representation is part of the mechanism for dealing with multiple-visit cases. Most of the decision rules (i.e. decision table columns) in IRIS require that something be true of the patient in a specified time period. Again, these specifiers are generally chosen from: past, present, past-or-present, future.

If an ISPEC is created using a create-expression that specifies a time of present, its TIME slot will contain the

date it was created on. It will therefore be able to satisfy decision rules that specify a time of present. At a subsequent visit, this ISPEC will no longer satisfy those decision rules. Instead, it will now satisfy ones that specify past-or-present or past.

### 3.5 Modifiers

The SB and TIME slots are necessary for all ISPECs. The SIDE slot is needed in a high percentage as well (relatively few nodes in IRIS's semantic net are SINGLE). There are other slots which we need only occasionally. For example, ISPECs for the node representing "intraocular pressure" need a VALUE slot, so that when we measure a patient's intraocular pressure, we can express that measurement with an ISPEC. Thus, an intraocular pressure of 35 would result in an ISPEC with a VALUE slot value of 35.

It would be wasteful to allocate a VALUE slot in all ISPECs, since only a few nodes require it in their ISPECs. DEGREE, COLOR, and WIDTH are other such occasionally-needed slots. We have dealt with this in IRIS by implementing a modifiers (MODS) slot in the ISPEC. The format of the value of the MODS slot is (property1 value1 property2 value2 ...). Thus, in the example above, the semantic net node "intraocular pressure" would get an ISPEC with a MODS slot value of (VALUE 35). This allows us to specify any slot we desire in any ISPEC without the overhead involved in explicitly defining that slot for all ISPECs.

We do not have a semantic net node representing "severely increased intraocular pressure". ISPECs representing severely increased intraocular pressure are actually ISPECs for the node "increased intraocular pressure" which have a MODS slot value of (DEGREE SEVERE). The MODS slot thus allows ISPECs to further specify the concepts represented by nodes.

If this were not the case, there would be a limit to the resolution of IRIS's representation for specific medical knowledge. If "increased intraocular pressure" were the highest resolution node available, then ISPECs without a MODS slot would provide us with no way of representing "severely increased intraocular pressure". The MODS slot thus allows us to achieve any level of resolution we desire in our representation of specific medical knowledge.

### 3.6 ISPEC Types

Each ISPEC has a TYPE slot. The value of this slot specifies the way in which its ISPEC is to be semantically interpreted. This is a feature which provides IRIS's knowledge representation

with considerable power. There is no limit to the number of ISPEC types that can be defined.

Currently, IRIS has seven ISPEC types. They are: NIL, FAMILYHISTORY, PTHISTORY, CHOSEN, COVEREDBY, SUBSUMEDBY, and TREATEDBY. NIL is the basic, standard ISPEC type. An ISPEC of type NIL means that the fact it represents is (or was) believed true (with the specified strength of belief) of the patient during the specified time. FAMILYHISTORY ISPECs represent facts which are true of people in the patient's family. PTHISTORY ISPECs represent facts about the patient which have been obtained from the patient.

ISPECs of type CHOSEN are very important in the operation of IRIS. These ISPECs are used to represent decisions made by the system. When IRIS chooses a diagnosis (by whatever strategy), the fact of this choice must be included in the store of patient-specific knowledge. IRIS will then be able to use this fact in its subsequent medical reasoning. The other ISPEC types represent the various sorts of inferences IRIS can make. A disease or treatment node is given a CHOSEN ISPEC when IRIS has decided to either make that diagnosis or perform that treatment.

SUBSUMEDBY ISPECs propagate from disease nodes which have enough positive strength of belief to be candidates for being CHOSEN. They propagate to other disease nodes which should not be CHOSEN when there is a sufficiently high level of belief in a subsuming disease.

When a disease node is given a CHOSEN ISPEC, it can then cause propagation of COVEREDBY ISPECs. COVEREDBY ISPECs propagate from CHOSEN disease nodes to symptom nodes that the disease can explain. This is part of IRIS's mechanism for knowing whether any symptoms remain unexplained.

When a treatment node is given a CHOSEN ISPEC, it can then cause propagation of TREATEDBY ISPECs. ISPECs propagate from CHOSEN treatment nodes to disease manifestations that the treatment has an effect on. This is part of IRIS's mechanism for knowing if every need for treatment can be satisfied by the current treatment regimen.

## 4. Decision Tables

IRIS's decision tables control propagation of ISPECs along the links of the semantic net. A decision table can be thought of functionally as a set of rules, but the actual mechanisms of a system based on decision tables differ somewhat from those of a rule based system.

A decision table in IRIS consists of a set of conditions, and a set of columns. Each condition is a predicate which tests whether the fact (or facts) it specifies is true of the patient. Columns function

as independent decision rules. Each column tests for a specified pattern of truth-values for a subset of the conditions. A column is said to be true if the truth-values of the conditions match its pattern. In this case, an ISPEC is created according to a specification also contained in the column. The ISPECs that propagate across a link are those created by the true columns of the link's decision table.

#### 4.1 Context Dependent Decision Tables

There are certain patterns of propagation which we need to specify for many different links in the semantic net. One such pattern is used for most of the "causes" links in the semantic net. The pattern can be paraphrased as: "If A causes B, and A is true, then conclude B with a strength of belief equal to the strength of belief in A times a constant associated with the causal link. Specify that B occurs in time after A."

We want to be able to implement such a pattern as a single decision table which can be associated with more than one semantic net link. This is accomplished in IRIS through the use of context dependent decision tables. Propagation is always done with respect to a particular link. That link is the propagation context. A context dependent decision table specifies how to create an ISPEC for the node on the far side of the link, if the node on the near side of the link is true. The identities of these nodes are obtained from the propagation context (the link being propagated with respect to), rather than from the decision table. Thus, we need to specify only one decision table that embodies the desired propagation pattern.

#### 5. The Propagation Process

The basic task facing any AIM system is to draw conclusions about the clinical situation, based on current knowledge of that situation. In IRIS, this task is performed by the propagation process. So far, we have described this process at the level of single semantic net links. While a description at this level can give an intuitive view of how propagation works, there are issues which become apparent only when propagation is viewed as occurring within a network, rather than along a single link. Propagation is a locally defined process which operates successfully because certain global aspects of its behavior have been taken into account.

The propagation process is designed to be quite general and flexible. We want a system which can automatically make all the desired inferences from a given set of facts with respect to a given model.

Furthermore, if we then change some fact or facts, we want the set of conclusions to be altered automatically to reflect the new set of facts. The results of the propagation process will then be dependent only on changes in the input data. Any such change will result in an appropriate change in the output of the model without any further intervention on the part of the user.

When propagation is attempted across a link, there is always the possibility that propagation has previously taken place across that same link. This complicates the logic of propagation somewhat. When propagation takes place across a link, the decision table associated with that link is evaluated. This produces a list of ISPECs for the target node. The target node may already have ISPECs which propagated from the source node, so the following three conditions must be checked for in determining whether to alter the ISPECs on the target node. If a column of the decision table which was formerly not true is now true, the ISPEC produced must be added to the target node's ISPECs. If a column of the decision table which was formerly true is not now true, the ISPEC it produced must now be deleted from the target node's ISPECs. Finally, if a formerly true column is still true, but has now produced an ISPEC which is different in some respect (e.g. the SB slot value is a function of the strength of belief in the node being propagated from), the ISPEC it produced must be replaced by the new one. These ISPEC matching tests are based on information carried within each ISPEC specifying the decision table and column which produced it. If any of the above conditions is true, the ISPECs of the target node are changed, leading in turn to propagation from the target node.

#### 6. Clinical Strategy

IRIS's clinical strategy is implemented through the six nodes "chosen diagnoses", "possible diagnoses", "unexplained symptoms", "chosen treatments", "possible treatments", and "untreated pathology". Any symptom for which an explanation is necessary propagates an ISPEC to "unexplained symptoms" when it is found to be true of the patient. In the same way, any condition which needs to be treated propagates an ISPEC to "untreated pathology". IRIS attempts to choose a diagnosis (or diagnoses) which will explain the maximum number of unexplained symptoms, and a treatment regimen which will halt as many pathological processes as possible.

Initially, IRIS collects information about the patient. At the end of this phase, IRIS's knowledge of the patient consists of the current set of ISPECs in

the semantic net. These ISPECs have been produced directly from answers to questions, or indirectly as the result of propagation. Any disease which has collected a sufficiently high CF via propagation will in turn have propagated an ISPEC to "possible diagnoses".

IRIS now chooses a diagnosis. Each disease node which has propagated an ISPEC to "possible diagnoses" is taken in turn and given an ISPEC of type CHOSEN. This causes COVEREDBY ISPECs to propagate from that disease node to all symptoms whose presence the disease can explain. Symptom nodes with COVEREDBY ISPECs do not propagate to "unexplained symptoms". At this point the number of symptom nodes still propagating to "unexplained symptoms" is noted. Comparing this to the number that propagated to "unexplained symptoms" before the disease was "chosen" gives us a measure of the disease's ability to explain the observed symptoms. The CHOSEN ISPEC is then removed from the disease node.

By doing this for each node that propagated to "possible diagnoses", we get a comparative measure of how the alternative diagnostic choices perform in explaining the observed symptoms. This measure of performance, along with the strength of belief which has propagated to the disease node are the criteria IRIS uses in choosing a diagnosis. Once the diagnosis node has been chosen, it is again given a CHOSEN ISPEC (permanently, this time). If there are still "unexplained symptoms", the process then repeats. At the end, some set of disease nodes will have been chosen, and each of these will have propagated an ISPEC to "chosen diagnoses".

The possible treatments for a disease are represented by nodes which are linked to the disease node by "treatmentfor" links. Propagation to the treatment nodes along these links takes place only when the disease node is chosen. Any treatment node which has accumulated a sufficiently high strength of belief in this way will propagate to "possible treatments" just as diseases do to "possible diagnoses".

IRIS decides on treatments using basically the same choice algorithm it uses for diagnostic decisions. Each treatment node which has propagated to "possible treatments" is in turn given a CHOSEN ISPEC. This causes TREATEDBY ISPECs to propagate from the treatment node to all nodes representing conditions this treatment can affect. They, upon receiving such ISPECs, cease propagating to "untreated pathology". From this process we get a measure of the alternative treatments' suitability for the patient. We then can choose one which gets a CHOSEN ISPEC permanently. Treatment nodes with CHOSEN ISPECs propagate to "chosen treatments".

We will illustrate this process in operation using a simplified (from a medical standpoint) example. In figure 6.1, the nodes "angle closure", "increased intraocular pressure", and "visual field loss" have received ISPECs as a result of user-supplied information. As a result, propagation (solid arrows) takes place from the three symptom nodes to "angle closure glaucoma", "unexplained symptoms", and "untreated pathology". Propagation then takes place from "angle closure glaucoma" to "possible diagnoses". In figure 6.2, "angle closure glaucoma" has received a CHOSEN ISPEC. Propagation to "possible diagnoses" ceases (broken arrow), and propagation now takes place to "chosen diagnoses" and "surgery". COVEREDBY ISPECs propagate from "angle closure glaucoma" to each of the three symptom nodes, "surgery" propagates to "possible treatments", and propagation from the three symptom nodes to "unexplained symptoms" ceases (a result of the COVEREDBY ISPECs). In figure 6.3, "surgery" is given a CHOSEN ISPEC, causing cessation of propagation to "possible treatments", propagation to "chosen treatments", and propagation of TREATEDBY ISPECs to the three symptom nodes. Propagation to "untreated pathology" from the three symptom nodes then ceases (caused by the TREATEDBY ISPECs).

This approach to clinical strategy enables us to restrict the points of contact between the semantic net/decision table knowledge base and the program that uses it. Only the six nodes mentioned at the beginning of this section are known to the program. IRIS's clinical strategy can be summarized as: 1) give CHOSEN ISPECs to nodes that propagate to "possible diagnoses" in such a way as to minimize propagation to "unexplained symptoms"; 2) give CHOSEN ISPECs to nodes that propagate to "possible treatments" in such a way as to minimize propagation to "untreated pathology"; 3) at the end of a session, produce output based on the ISPECs of nodes that are currently propagating to "chosen diagnoses" and "chosen treatments".

The advantage of such a restricted area of contact between the program and its knowledge base is that the knowledge base can be extensively altered without requiring changes in the program. The program only requires of its knowledge base that the six nodes be present, and that giving CHOSEN ISPECs to disease and treatment nodes will have effects on propagation to "unexplained symptoms" and "untreated pathology". The exact mechanism of these effects is unknown to and irrelevant to the program. It need not be done via COVEREDBY and TREATEDBY ISPECs.

This research was supported in part by grant RR-243 of the NIH Biotechnology Resources Program.

### References

Brown, J.S., Bell A.G., Burton R., Sophisticated Instructional Environment for Teaching Electronic Troubleshooting, Report AFHRL-TR-74-77, Technical Training Division, Lowry AFB, Colorado, 1974.

Kulikowski, C.A., Safir, A., Trigoboff, M., and Weiss, S.M., "Clinical Consultation and the Representation of Disease Processes: Some Artificial Intelligence Approaches", in Proceedings of the 1976 AISB Conference, Edinburgh, Scotland, 1976.

Kulikowski, C.A. and Weiss S.M., Computer-Based Models for Glaucoma, Report CBM-TR-3, Department of Computer Science, Rutgers U., 1971.

Pauker, S.G., Gorry, G.A., Kassirer, J.P., and Schwartz, W.B., "Towards the Simulation of Clinical Cognition", The American Journal of Medicine, 60:981, June 1976.

Pople, H.E. Jr., Myers, J.D., and Miller, R.A., "DIALOG: A Model of Diagnostic Logic for Internal Medicine", in Proc. 4th IJCAI, 2:849, 1975.

Rubin, A.D., Hypothesis Formation and Evaluation in Medical Diagnosis, MIT Artificial Intelligence Laboratory Technical Report AI-TR-316, January 1975.

Shortliffe, E.H., MYCIN - A Rule-Based Computer Program for Advising Physicians on Antimicrobial Therapy Selection, AI Memo 251, Computer Science Department, Stanford U., 1974.

Trigoboff, M., Ph.D. thesis, in preparation, 1977.

Trigoboff, M., "Propagation of Information in a Semantic Net", in Proceedings of the 1976 AISB Conference, Edinburgh, Scotland, 1976.

Weiss, S.M., A System for Model-Based Computer Aided Diagnosis and Therapy, Report CBM-TR-27, Department of Computer Science, Rutgers U., 1974.

Woods, W.A., "What's in a Link: Foundations for Semantic Networks", in Bobrow, D. and Collins, A. (eds), Representation and Understanding, Academic Press, New York, 1975.

