NSF Director's Program Review - 1976

by

Sally Sedelow, Program Founder & Director, Intelligent Systems

This presentation for the Intelligent Systems program is going to proceed somewhat heuristically--through parallel and analogy--in an effort to show, first, the current interconnectedness of approaches to computer-based intelligent systems with approaches to the study and description of human intelligent systems, especially their language capabilities, second, the interrelatedness between different areas of research on computer-based intelligent systems--areas such as scene understanding and man-machine dialogue in natural language--which might on the surface seem quite separate from each other, and third, some research questions deriving from the need for intelligent systems to employ languages. The stress upon interconnectedness and interrelatedness should not be taken as in any way a contradiction of the emphasis on the long-range plan for the Intelligent Systems program upon complementarity between man and machine. We are looking toward the use of the computer to perform tasks which complement human performance--many of those tasks, such as pre-sorting masses of data or providing otherwise unavailable information through dialogue with the human being, entail a symbolic processing capability on the part of the computer. Hence, the inter-relationships explored here are central to the achievement of many of the desired complementarities.

Preliminarily, it should be noted that the centrality of a language capability for all forms of intelligent systems (both human and non-human) is taken to be now self-evident. By languages are meant natural languages, programming languages, graphics languages, indeed any symbolic representations and mediations of perception and analyses of perceptions. The various languages used in intelligent systems would seem to be interrelated. At any rate, the kinds of notational schemes and structures (which, of course, are themselves languages) suggested for representing the knowledge a computer system has about the world show interesting parallels both among themselves and with representations suggested for describing human natural language capacity. Such interrelationships are being explored by research scientists from a variety of disciplines. This presentation will allude to some of that research, especially where that research is directed toward the languages such as English which are, or at least have been tacitly assumed to be most 'natural' to human intelligent systems. For the immediate future, at least, it is with human intelligent systems that artificially intelligent systems will most frequently interact, thus, it is in these so-called "natural languages" that human intelligent systems and computer-based intelligent systems will frequently find it convenient to communicate.

I want to turn now to the first main topic of this presentation--the interconnectedness of current approaches to computer-based intelligent systems with approaches to the study and description of human intelligent systems, especially their language capability. Parallels will be drawn between two major paradigms--the first, an aspect of the transformational paradigm for natural language representation, and the second, the frame paradigm for representing knowledge in computer-based systems. [redacted] The two paradigms are among the most used in their respective fields and the parallels between them are sufficiently striking as to demonstrate the interconnectedness mentioned earlier. What if anything does the interconnectedness signify? That topic will be discussed more fully later, but at the least the interconnectedness does represent a preserving of continuity in symbol representation across two forms of intelligent systems--a continuity which may help provide some useful insights into how to set up computer-based systems so as to deal more adequately with symbolic information.

Now to the example: displayed on slides will be the so-called phrase-structure component of one of Noam Chomsky's earliest attempts (Syntactic Structures, 1957) to describe natural language within the transformational paradigm. While the slides are being projected, I want for comparison to read aloud parts of Marvin Minsky's description of the "frame," set forth in his paper "A Framwork for Representing Knowledge." Minsky says that

> A frame is a data-structure for representing a stereotyped situation...We can think of a frame as a network of nodes and relations. The "top levels" of a frame are fixed, and represent things that are always true about the supposed situation."

Now notice for comparison that in the transformational diagram the top level is occupied by an entity which in Chomsky's view of language is also fixed -- that is, the sentence. Lower levels in Chomskys' representation are occupied by entities which, as the level descends, grow more specific to a particular sentence until the actual words in a sentence, called the terminal string, appear at the lowest level. Now, to return to Minsky, after noting that "the top levels of a frame are fixed, and represent things that are always true about the supposed situation," he goes on to say that "the lower levels have many terminals - 'slots' that must be filled by specific instances or data."

As the next slides show, implicit in Chomsky's tree are procedures, in this case rewriting rules (represented by arrows), indicating how one can begin with an abstract sentence and terminate with "the man hit the ball." Like Chomsky's tree, Minsky's frame (for which he provides no diagrammatic description) contains both data and procedures, although the representation of knowledge in his frame is considerably more complicated than Chomsky's early description of the phrase-structure component of a transformational grammar. But it should be noted that Chomsky's representations have themselves grown ever more complex as he packs more specific data information into his representation of knowledge about natural language. The parallels between these two representations ought not to be overstressed but they do seem heuristically interesting when considering modes of representing information and actions central to intelligent systems.

Another parallel, suggestive of interconnectedness, between linguists' approaches to natural language and computer scientists' approaches to the development of intelligent systems, seems worth mentioning. That is, until very recently within the last half-century, linguists thought it appropriate to divide form from content and to emphasize form -- for example, syntax -- at the expense of content -- for example, semantics. Early research on computer-based intelligent systems also tended to emphasize form, embodied in problem-solving procedures, and, insofar as possible, to ignore knowledge or content. Thus, relatively early, computer scientists developed a General Problem Solver which couldn't solve many problems because essential interrelationships between problem-solving methods and problem-content were ignored. Now, a number of linguists assert that semantics has primacy over syntax and many computer scientists argue for a strong infusion of knowledge detail into any problem-solving system. This shift accounts for the pre-occupation of some linguists with the representation of semantic knowledge and the preoccupation of some computer scientists with the representation of knowledge, a representation which is often termed a semantic representation.

One other specific interrelationship between the languages of computer-based systems and human intelligent systems should be mentioned: these the types of formal languages which have provided the basis for formal mathematical study and development of aspects of programming languages as well as their translators, such as compilers, are directly borrowed from the efforts -- notably by Chomsky -- to provide formal approaches to the description of natural languages. That is, the transformational grammars used for the generalization and study of natural languages are the Type Zero languages shown in the slide, so called context-sensitive grammars are Type One grammars as shown in the slide, context-free grammars Type Two, and so on. It is apropos to note that these formal grammars can be represented by types of automata, or abstract machines. Thus, a central theoretical core of computer science can be shown equivalent to formal descriptions of languages derived directly from work on the languages natural to human intelligent systems. Again, while as ways of making computer-based systems behave more intelligently are explored we have found an interrelationship which seems heuristically significant. Indeed an aspect of this interrelationship has already been exploited, through the design of computer-based natural language parsers based on the notion of transitions from one state of an automaton to another state of that automaton.

Now, I would like to turn briefly to the second major topic of this presentation, the interrelatedness between areas of research on computer-based intelligent systems that might on the surface seem quite separate from each other.

The "frame" is proposed as a structure for representing not only procedural and factual knowledge but also for a range of types of knowledge, including visual perceptions as well as linguistic descriptions of all types of perceptions. The effort to find comprehensive frameworks for representing knowledge is of particular interest in intelligent systems research because a number of tasks for which intelligent systems might seem desirable simultaneously entail different types of perceptions and responses. Thus knowledge representation is correctly seen as a core area of research in intelligent systems.

One research effort (at the University of Pennsylvania) currently supported by the Intelligent Systems Program is endeavoring to borrow a representation scheme used for simple two-dimensional shapes and put it to new use, representing natural language data. The recognition of shapes is central to the area of scene understanding, which forms one major strand of current research on intelligent systems. The representation of natural language data is of interest to the Pennsylvania group because their research is directed toward facilitating man-machine dialogue in English. The next two Figures (8 & 9) show the lattice representation used both for the shape recognizer and for the lexical string recognizer. The general goal, for both the shape recognizer and the lexical string recognizer, is to succeed in matching a pattern when only partial information is available. Thus, in the case of the Shape Recognizer in Figure 8, Node A is identified even when only P1 and P3 are available. P2 is "covered" by the arc extending from A to P3. In the case of Figure 9, if inputs contained the notions of "Benign" and "Give," then the notion of "Belong" would also be directly accessible under Schemata-1 (Borrows) even though this latter notion does not appear in the input. The significant point here is not the detail contained in the lattices but rather, at a more abstract level, the fact that the lattice structures which seem suitable for the representation of shapes also seem useful for the representation of natural language strings.

Another issue raised in the research at Penn leads directly into this presentation's third major topic, which is identification of some of the research issues deriving from the need for intelligent systems to employ languages. The issue raised is whether the approach to knowledge representation should be primarily procedural or declarative. This issue is by no means restricted to language research, in one sense, but since knowledge of all forms can almost always be said to be represented by language (and on digital computers it is), in a larger sense language is the point at issue here. The research group at Penn favors a declarative approach to the representation of knowledge: this approach separates the data from the procedures which manipulate it. They argue that by factoring out the procedural component and making the data structure a precisely defined entity, increased complexity can be managed.

One of the early prime advocates (Winograd) of the procedural approach, which embodies the facts or information in the procedures which manipulate that information, is also supported by the Intelligent Systems Program. In a recent paper entitled "Frame Representations and the Declarative/Procedural Controversy," Winograd concludes that what is needed is a synthesis of the two approaches. He contends that the flexibility and economy of declarative knowledge come from the ability to decompose knowledge into "what" and "how". Learnability and understandability of declarative knowledge are said to come from the strong independence of the individual axioms or facts. On the other hand, Winograd believes that procedures give an immediate way of formulating the interactions between the static knowledge and the reasoning process, and allow a much richer and more powerful interaction between the "chunks" into which knowledge is divided. Winograd argues that in trying to achieve a synthesis we must ask not "How can we combine programs and facts?", but "How can our formalism take advantage of decomposability without sacrificing the possibilities for interaction?" As a step toward such synthesis, Winograd and Daniel Bobrow are designing a Knowledge Representation Language intended to facilitate both declarative description and procedural description.

On the subject of declarative description, Winograd and Bobrow believe "that the description of a complex object or event cannot be broken down into a single set of primitives, but must be expressed through multiple views." In contrast to logical notation, they believe "it is more useful and perspicuous to preserve in the notation many of the conceptual differences which are reflected in natural language, even though they could be reduced to a smaller basis set." They hold "that multiple descriptions containing redundant information are used in the human representation system to trade off memory space for computation depth, and that computer systems can take advantage of the same techniques." They note that the choice of where to put redundancy provides further structure for memory and can be used to limit search and deduction. As a simple example, they suggest an understander system which might know that every plumber is a person and that Mary is a plumber. The memory unit for Mary would contain a descriptor stating that she is a plumber and would very likely also contain an explicit descriptor saying that she is a person. This is redundant, but without it the system would be continually rededucing simple facts, since "personhood" is a basic property often used in reasoning about entities. Memory structure in KRL (Knowledge

Representation Language) is organized in a way which makes it possible to include redundant information for immediacy while keeping the ability to derive information which has not been explicitly stated.

Winograd and Bobrow's approach is perhaps best summarized in a statement by Winograd elsewhere: "Current research and that of the next few years will be strongly focused on exploring different representations, and in particular the effect they have on trying to build large knowledge bases. Many elegant schemes seem to break down when we try to include a domain whose size realistically approximates even a small sub-part of a person's knowledge. Our explorations of representation must pay special attention to problems of modularity and interconnection which will make large systems possible, and must underly the human reasoning system. By focusing on what is common in intelligent systems, we are working toward a general theory of representation and knowledge."

With reference specifically to natural language, Winograd notes that "an utterance is like a rough sketch which allows the hearer to reconstruct the meaning." Another research effort supported by Intelligent Systems, this one directed toward building computer-based consultants, is attempting to deal directly with some of the implications of the "rough sketch" nature of an utterance. Examples drawn from this research, being conducted at the Stanford Research Institute, will deal with the finest level of detail under Topic 3 of this presentation: research questions deriving from the need for intelligent systems to employ languages. Although the examples may look trivial, the problems they exemplify are not at all trivial in the context of understanding natural languages. In part, trivial-appearing examples are used so as to underline the complications posed by even simple-looking problems. For example, one of the problems to which the SRI group is addressing itself is the determination of antecedence of definite noun phrases. Knowing the antecedence of noun phrases introduced by the definite article is knowledge any natural language understanding system should have. But such knowledge is not as easily come by as might be supposed. One approach to what would seem to be a straightforward problem has been to keep a history of the dialogue and to have the computer go back as far as necessary to find the object mentioned. A problem with this approach is that the system may remember too much; in particular, it may consider some item as still in focus when the user has forgotten it. For example, given the sentence: "Is the wrench lying on the table broken?" (Figure 11), the system may remember some wrench mentioned earlier in the conversation and assume that to be the wrench in question, concluding that "is lying on the table" is the verb phrase. In fact, as the human intelligent system can see, "lying on the table" modifies "wrench" and the wrench "lying on the table" may or may not be the same as some wrench mentioned earlier in the conversation. How does the computer system know which "wrench" is in question? Or, to take an example illustrating another problem, we have the sentences shown in ~~slide 11~~ Figure 12. In this case, the system has to know that the cover is going to be attached with screws. Such knowledge must somehow be represented in the system's memory so that the reference for, "the screws," can be ascertained. The SRI group proposes to supply the system with models of tasks being discussed which will enable the system to deduce the relevance of, for example, "the screws." An example of an overall task model for a pump installation is shown in Figure 13. Presumably, a much elaborated set of boxes for the one in this figure stating "Connect Belt Housing Cover to Belt Housing Frame" would include information about screws being used to do the connecting.

The SRI research explores a range of approaches to problem-solving as well as to questions of general interest to the study of natural languages -- topics such as fragmentary and anomalous sentences and the treatment of time. Within these larger formal contexts, it is also necessary to cope with content. The task model, or plot, shown in Figure 13 is one of several current approaches to the representation of knowledge specific to a given task. As should be evident by this point, a cutting edge in Intelligent Systems research entails the interfaces between form and content and between procedures and declarative knowledge: What should the balance be? What should the form of the interface be? How should the procedures and knowledge be represented?

There is not time to refer to the many interesting problems to be explored under the rubric of Intelligent Systems, or even under the more restricted aspects of intelligent systems discussed in this presentation; (See Figure 14). The goal in this presentation has been to convey through analogies and parallels, first, the current interconnectedness of computer-based approaches to intelligent systems with approaches to the study and description of human intelligent systems, especially as to their language capabilities, second, the interrelatedness between areas of research on computer-based intelligent systems--areas such as scene understanding (shape recognizer) and man-machine dialogue in natural language--which might on the surface seem quite separate from each other, and, third, some research issues deriving from the need for intelligent systems to employ languages.

It is surely legitimate to ask whether the perceived interconnectedness is no more than an accident of paradigmatic fashion; that is, certain approaches come and go and, when popular, an effort is made to apply them to 'everything.' The answer to this question is that while, of course, approaches do come and go, some do prove to be enormously productive of scientific insight. In the case of intelligent systems, an interesting point at the moment is that so many different modes of perception and structurings of perception are being bound up together through representation schemes which are parallel. This binding together may ultimately prove to have been overly attractive, but it also could prove a source of strength as attempts are made to define and integrate the various capacities characteristic of intelligent systems, whether biological or artifactual.

Slides

OVERVIEW

I. COMPUTER-BASED INTELLIGENT SYSTEMS ⟷ HUMAN INTELLIGENT SYSTEMS (LANGUAGE CAPABILITIES)

II. SCENE UNDERSTANDING ⟷ MAN-MACHINE DIALOGUE

III. RESEARCH QUESTIONS DERIVING FROM INTELLIGENT SYSTEMS' NEED FOR LANGUAGES

MCS77-195
11-16-76

Figure 1

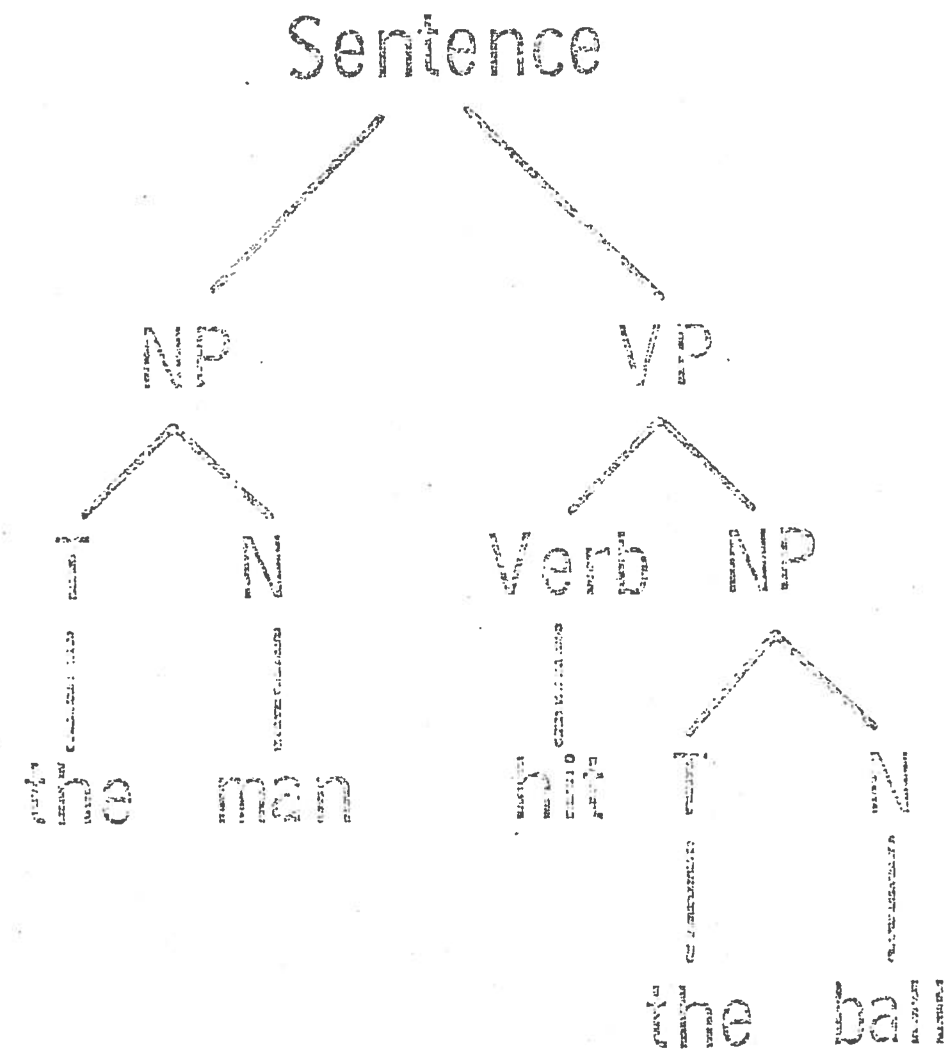

MCS77-101
11-16-76

13 (i) Sentence → NP + VP

(ii) NP → T + N

(iii) VP → Verb + NP

(iv) T → the

(v) N → man, ball, etc

(vi) Verb → hit, took, etc.

MCS77-102
11-16-76

Sentence

NP + VP (i)

T + N + VP (ii)

T + N + Verb + NP (iii)

the + N + Verb + NP (iv)

the + man + Verb + NP (v)

the + man + hit + NP (vi)

the + man + hit + T + N (iii)

the + man + hit + the + N (iv)

the + man + hit + the + ball (v)

MCS77-103
11-16-76

Formal Languages	Symbolic Representations	Grammars
Type 0	$\alpha \longrightarrow \beta$	Transformational
Type 1	$\alpha_1 A \alpha_2 \longrightarrow \alpha_1 \beta \alpha_2$	Context-Sensitive
Type 2	$A \longrightarrow \beta$	Context-Free
Type 3	$A \longrightarrow aB$ or $A \longrightarrow a$	Regular or Finite State

MCS77-104
11-16-76

EQUIVALENCES OF AUTOMATA AND FORMAL LANGUAGES

AUTOMATA	FORMAL LANGUAGES
Turing Machines	Type 0
Linear Bounded Automata	Type 1
Pushdown Automata	Type 2
Finite Automata	Type 3

MCS77-107
11-16-76

GRAMMARS	FORMAL LANGUAGES	AUTOMATA
Transformational	Type 0	Turing Machines
Context-Sensitive	Type 1	Linear Bounded Automata
Context-Free	Type 2	Pushdown Automata
Regular or Finite State	Type 3	Finite Automata

MCS77-105
11-16-76

SHAPE RECOGNIZER LATTICE
(PARTIAL INFORMATION)

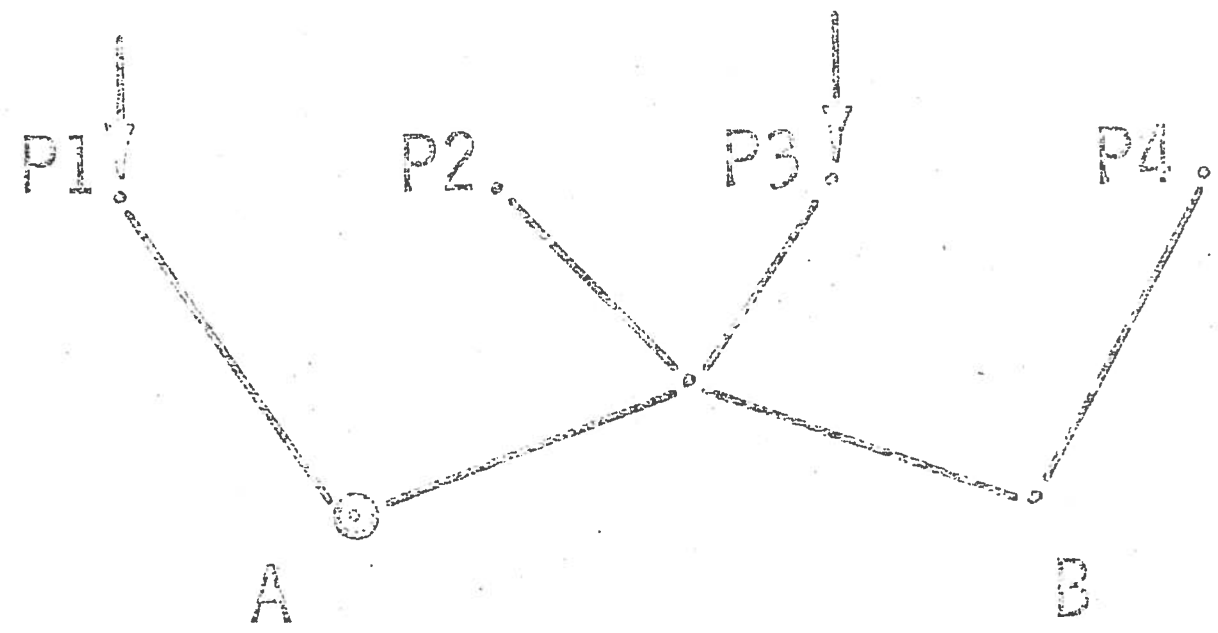

Figure 8

MCS77-103
11-16-76

LEXICAL STRING RECOGNIZER LATTICE

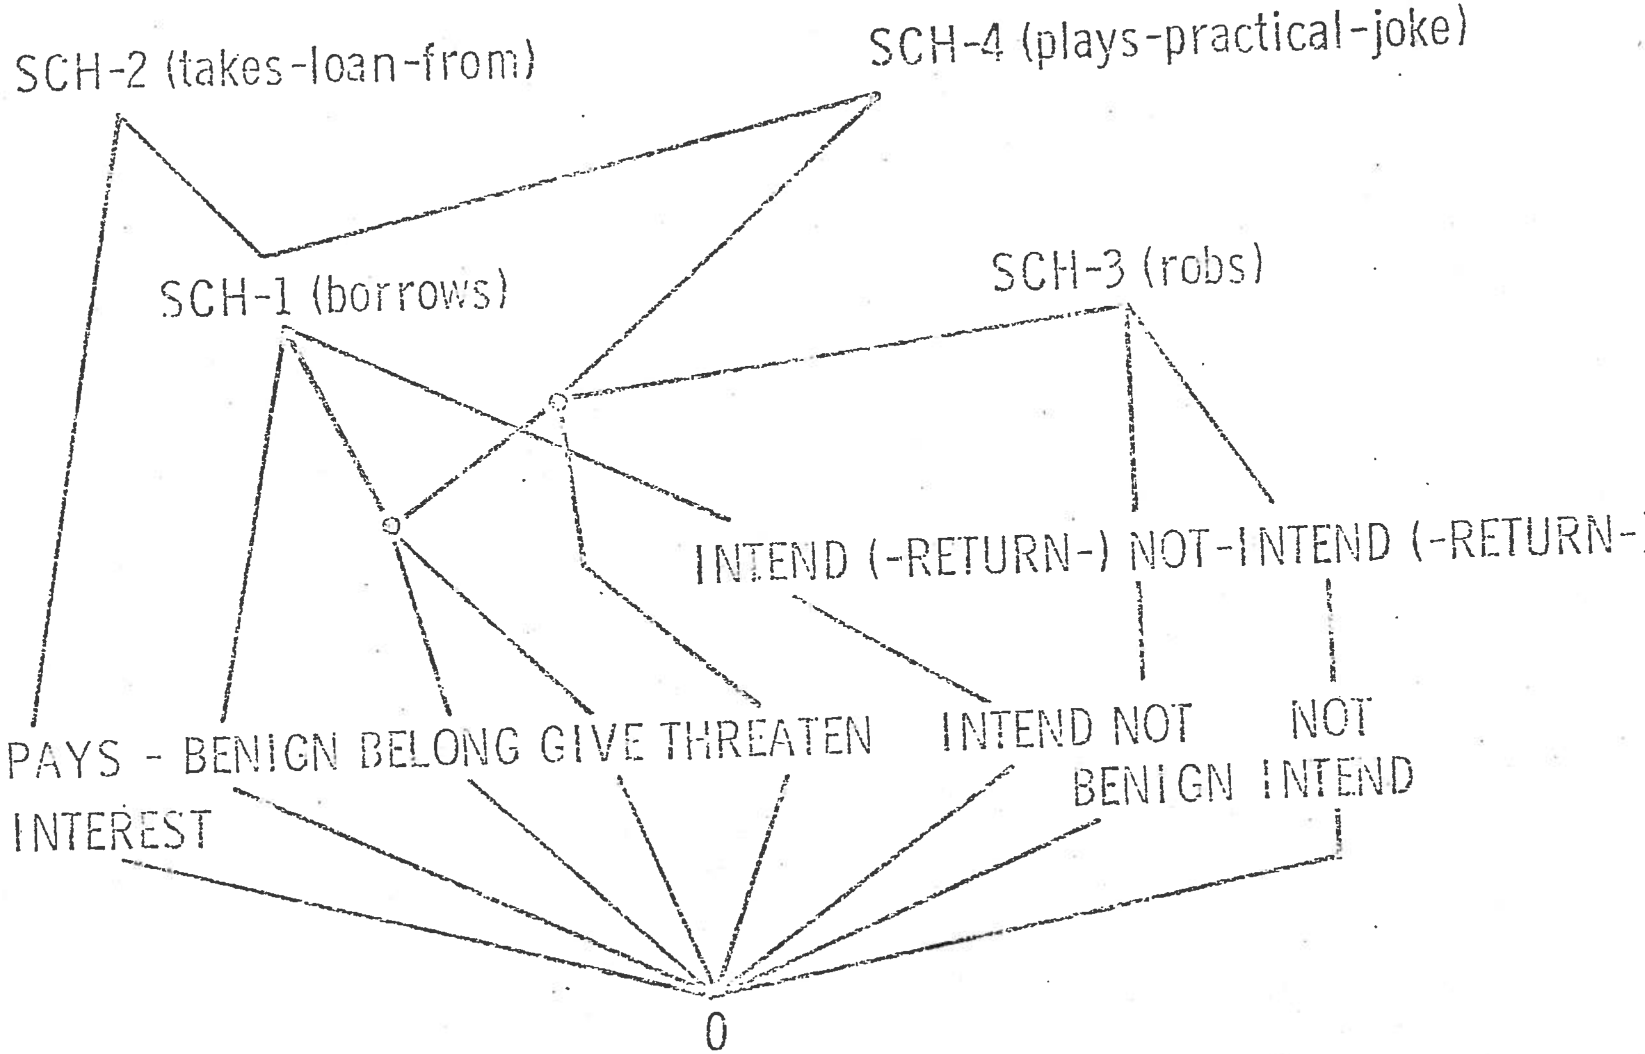

Figure 9

MCS77-103
11-16-76
32

Declarative Approach

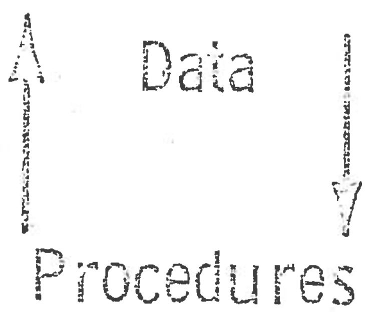

Procedural Approach

Procedure 1

Data

Procedure 1

Data

Procedure 2

Data

Procedure 1

MCS77-109
11-16-76

IS THE WRENCH LYING ON THE TABLE BROKEN?

Figure 11

MCS77-110
11-16-76
35

ATTACH THE COVER TO THE CONTAINER.

THE SCREWS ARE IN THE BOX ON THE TABLE.

Figure 12

MCS77-111
11-16-76

PUMP INSTALLATION

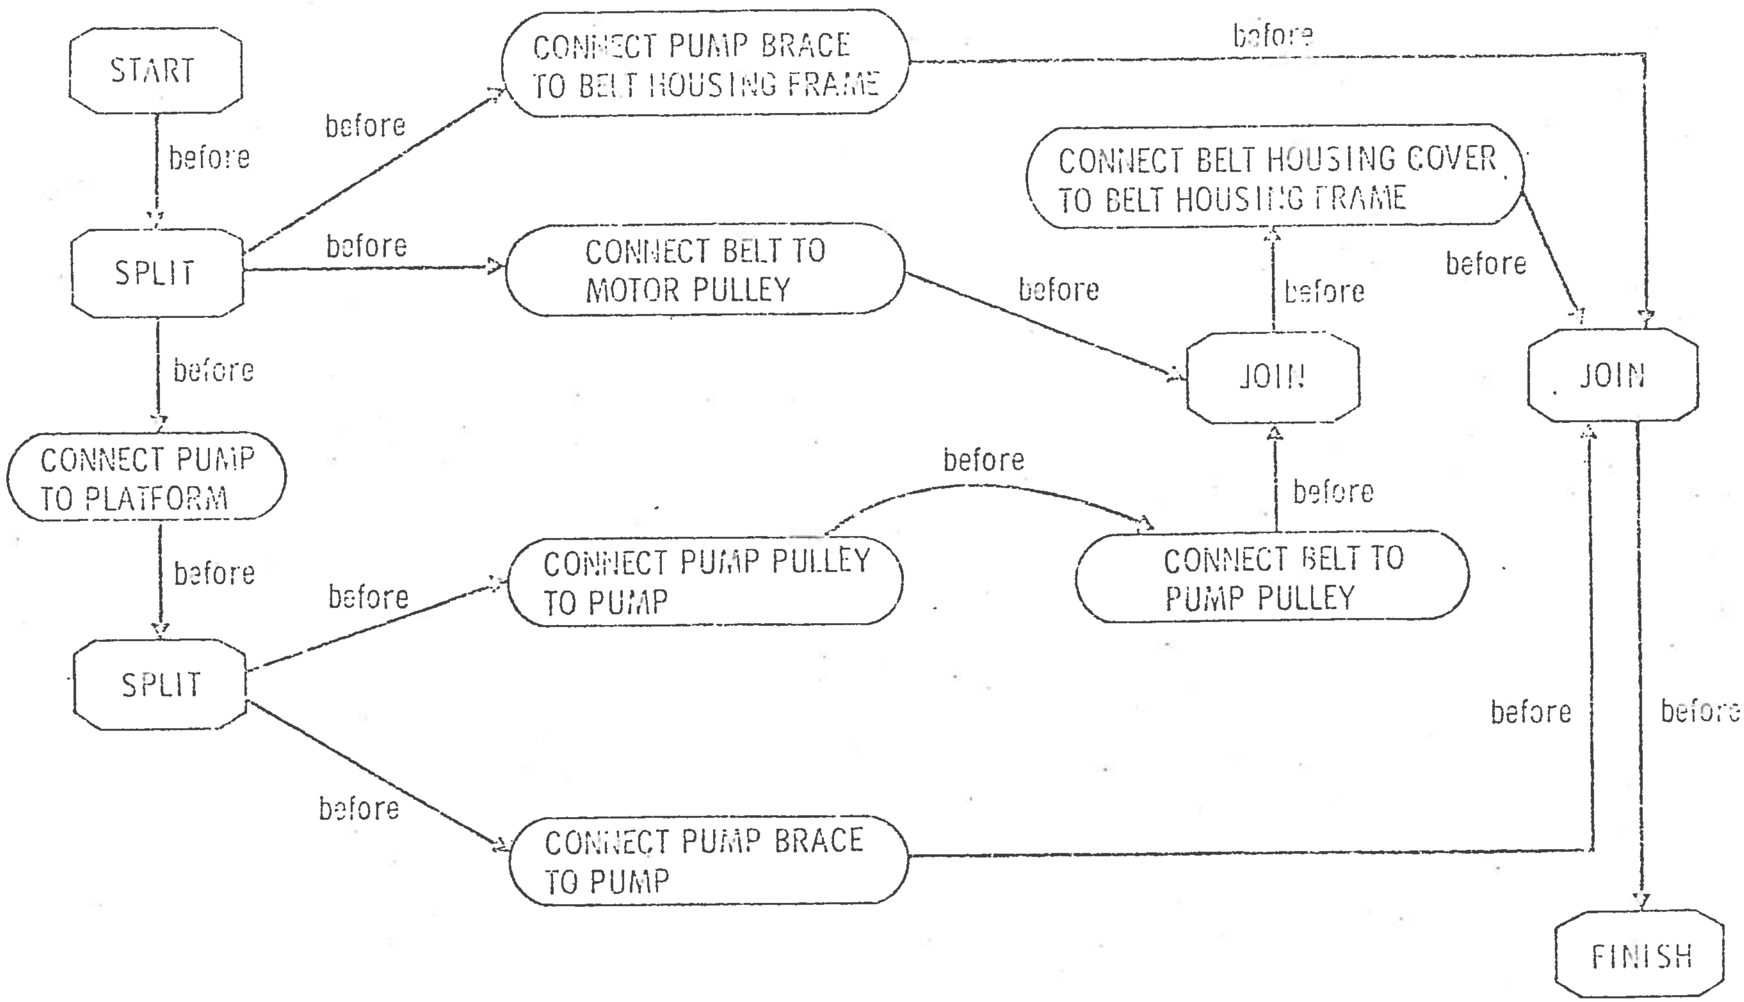

Figure 13

MCS77-112
11-16-76

SUMMARY

I. COMPUTER-BASED INTELLIGENT SYSTEMS ⟷ HUMAN INTELLIGENT SYSTEMS (LANGUAGE CAPABILITIES)

II. SCENE UNDERSTANDING ⟷ MAN-MACHINE DIALOGUE

III. KNOWLEDGE REPRESENTATION

A. DECLARATIVE -- PROCEDURAL

B. CONTENT -- FORM

1. TASK MODELS

2. PARTIAL INFORMATION

3. TOPICS SPECIFIC TO NATURAL LANGUAGE, e.g. DEFINITE NOUN ANTECEDENCE

MCS77-137
11-16-76

Figure 14