

NOT FOR QUOTATION
WITHOUT PERMISSION
OF THE AUTHOR

**EXPERT VS MANAGEMENT SUPPORT SYSTEMS:
SEMANTIC ISSUES**

Ronald M. Lee

February 1983
WP-83-37

Working Papers are interim reports on work of the International Institute for Applied Systems Analysis and have received only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute or of its National Member Organizations.

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
2361 Laxenburg, Austria

ABSTRACT

Expert systems hold great promise for technical application areas such as medical diagnosis or engineering design. They are, we argue, less promising for management applications. The reason is that managers are not experts in the sense of possessing a formal body of knowledge which they apply. The limitations of artificial intelligence approaches in managerial domains is explained in terms of semantic change, motivating attention towards management (decision) support systems.

Keywords: expert systems, management decision support systems, knowledge representation, formal semantics, applied epistemology.

CONTENTS

INTRODUCTION	1
BUREAUCRACIES VS MACHINES	4
ORGANIZATIONAL ADAPTATION	8
MECHANICAL VS MANAGERIAL COGNITION	10
A. The Ego Problem	11
B. The Semantic Problem	14
EXPERT SYSTEMS VS DECISION SUPPORT SYSTEMS	21
SUMMARY REMARKS	23
POST SCRIPT	25
ACKNOWLEDGMENTS	25
REFERENCES	27

**EXPERT VS MANAGEMENT SUPPORT SYSTEMS:
SEMANTIC ISSUES**

by

Ronald M. Lee

INTRODUCTION

AI is getting market appeal. Expert systems, robotics and 5th generation technology are getting serious recognition in the economic plans for 1984 and beyond. The attempt here is to assess the potential impact of AI future technology on commercial organizations and other social institutions. Technology assessment suffers the lack of a convincing methodology. Hence the strategy here is not to try to predict the actual course of AI innovations, but rather consider what would be the theoretical limits to the technology.

Our concern is mainly with AI technology in *organizations*, i.e., with groups of people working in cooperation. These remarks are not intended to apply to industrial robots, nor to single user expert systems, but rather to what might be called a 'knowledge-based information system'

(KBIS). Such applications would seem to be the eventual result of a convergence of database management with AI knowledge representation.

To simplify the argument and avoid large literature surveys, we take the liberty of imagining a future KBIS as a large scale theorem prover operating on a database of logical assertions about the organization and its environment. This trend might be discerned from the literature on 'logic and databases' (Gallaire et al. 1978, 1981) and the logic programming discussions of relational databases (Clocksin and Mellish 1981, Coelho 1980, Kowalski 1979b).

The question is, what could such a KBIS do?

The principle function of an information system in organizations is to facilitate communication between individuals that are geographically and/or temporally separated. Unlike e.g., telephone or electronic mail, the advantage offered by an information system accessing a structured database is that it offers the possibility of making *inferences* on the communications it intermediates. Inferencing facilitates the chunking of information (Miller 1956) necessary as communications flow upward in the management hierarchy (Jacques 1976).

Jay Galbraith (1973, 1977) observes that hierarchy itself is an information processing device, helping the organization to cope with the conflicting pulls of a complex environment vs the limited attention and bounded rationality of management (Simon 1955). Knowledge-based information systems would, we expect, reduce the complexity by taking over more and more managerial problem solving.

But is there a limit? Wouldn't the future, super-powerful, knowledge-based information system eventually eliminate the need for management? The arguments which follow lead to a negative conclusion. AI will make an important contribution to management problems, but the brave new world of the future will not *only* be built with technology.

Arguments of this sort tend to rely on the 'unstructuredness' of the managerial task as the basis for a view that information technology will at best aid, but not replace management (Gorry and Scott-Morton 1971, Keen and Scott-Morton 1978). But that argument eventually encounters a circularity if by 'structured problems' is meant those that have a decision algorithm. Technology has an untidy habit of advancing beyond problems that were previously thought impossible.

The arguments given here are based on two interconnected themes. One is the problem of preferences (goals, values, free will), which we argue that computers don't have. (Computers don't intrinsically prefer chocolate to vanilla.) The other theme involves basic issues in semantics which, especially for organizations in dynamic, uncertain environments, provide fatal difficulties for even an idealized AI system.

The arguments, interestingly, have a certain parallel with issues of bureaucracy. Various insights can perhaps be exchanged between AI knowledge representation topics and the apparent limitations to bureaucratic rationalization.

BUREAUCRACIES VS MACHINES

A characteristic of machine intelligence is that it is 'rule based'. If we consider only this software aspect (and ignore differences in processor hardware), then the most ubiquitous and successful examples of mechanical cognition are bureaucracies. Yet while the projects to create various types of artificial intelligence have a certain romance and intellectual adventure about them, the term 'bureaucracy' seems at best dreary and more often spiteful. It is laden with negative connotations of plodding, brutish organizations, insensitive to the individual; indifferent to the exceptional.

Yet in more scientific usage, 'bureaucracy' is used neutrally as merely one form of administration. The negative associations it has in popular usage gives empirical evidence that people's encounters with bureaucracies are often unpleasant. The definition of bureaucracy used here is based on Weber (1956/1978), indicating organizations whose administration is based on explicit rules and procedures. This contrasts with an idiosyncratic form of management based on personal interest and the whims of the moment. Bureaucracies, then, are organizations whose behavior is 'rationalized' to eliminate such idiosyncratic tendencies. This gives rise to a concept of organizational *role*, and explicit, detailed job descriptions. Personnel become substitutable; the organization takes on a mechanical consistency and permanence that outlives its members. In Weber's words,

Bureaucracy develops the more perfectly, the more it is "dehumanized," the more completely it succeeds in eliminating from official business love, hatred, and all purely personal, irrational, and emotional elements which escape calculation (Weber

1956/1975:975).

Consider how this view compares with standard models of computation. In automata theory (e.g., Hopcroft and Ullman 1969), we view a computer abstractly as a language processor, transforming an input string of symbols to output symbols (see Figure 1a). In information systems applications we can regard these symbols as part of a common language, call it L_{RW} , which are assertions about the 'real world' (organizational environment). These assertions are normally stored in the organization's database and the processor is invoked by queries, calls to application programs, etc. Hence, what we call the 'automaton' here is meant to include the entire set of application programs, DBMS software, query interfaces, etc. (in whatever future software designs you like).

The automaton, as language processor, is regarded as a grammar. This grammar is itself defined in a notation, call it L_C . Practically, L_C corresponds to an arbitrary programming language.* Ignoring efficiency considerations, we might regard L_C as reducing to a set of production rules (Davis and King 1975) of the form

IF <condition> THEN DO <action>.

If none of the various conditions are met, that is, if no rule is actuated, the default is inaction. The machine doesn't do anything it's not instructed to do by one of its rules.

* It is common for LISP users and others to deny the distinction between data and program. Distinctions however depend on expository purpose. We could of course consider the language formed in the union of L_{RW} and L_C . The two languages are distinguished semantically. The semantics of L_{RW} is all those expressions in the information system which denote real world phenomena. The semantics of L_C is machine operations. These are of course hopelessly intertwined in all present day implementations, which is why we resort to talking about idealized machines.

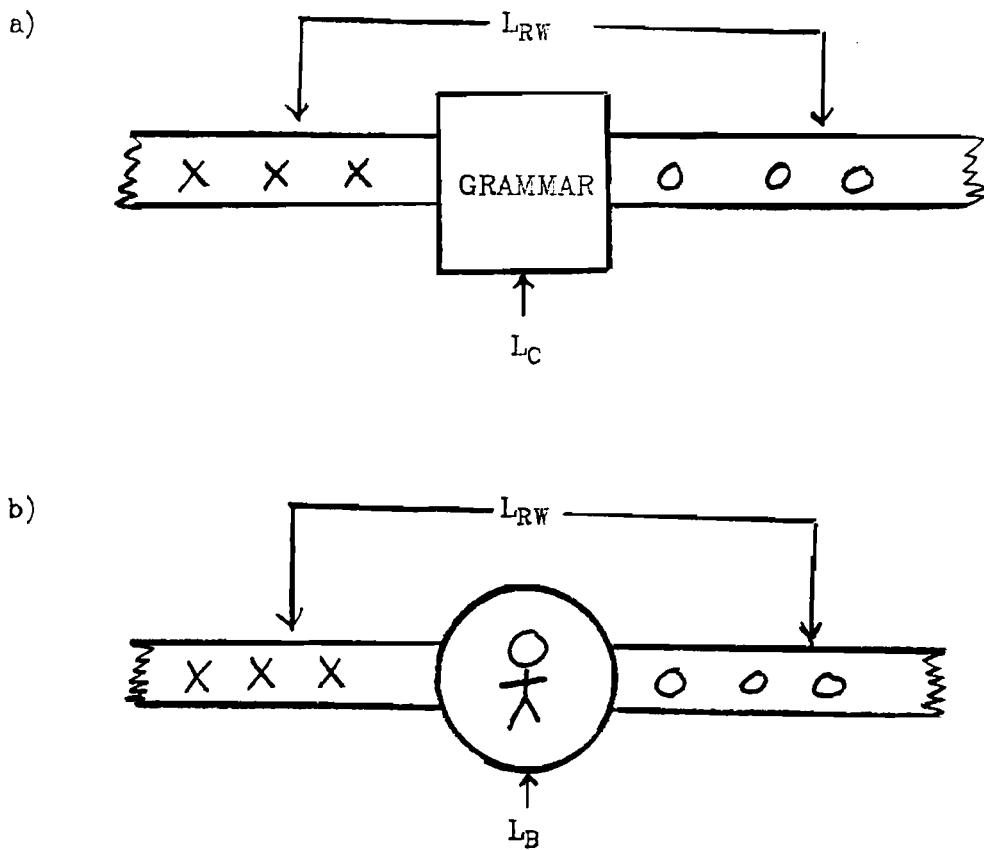


Figure 1.

A currently popular view of organizational management (e.g., March and Simon 1958) regards managers as information processors. Taking the metaphor literally, we might replace the automaton with a person (Figure 1b). The 'programming' of this person might be in another language, L_B , expressing the various bureaucratic rules and procedures this person is to follow.

But if we regard L_B (bureaucratic programming) abstractly in the way we did L_C (computer programming), we encounter a problem if we use only production rules. As observed in a body of literature in

organizational psychology and sociology (e.g., Maslow 1943, McGregor 1960, Cyert and March 1963, March and Olsen 1979) people are not naturally idle. They have their own individual interests, goals, aspirations, etc. which they are seeking to satisfy through their participation in the organization.

When these correspond to the interests and goals of the organization itself, we tend to regard their independent behavior as 'initiative', otherwise it is considered more as the dysfunctional pursuit of 'personal interest'. L_B (bureaucratic programming) therefore contains another basic aspect. It not only orders the execution of desired behavior, but *restrains* the performance of undesired behavior. In Lee (1980), we suggest that a primitive structure of bureaucratic software would therefore include the basic operators of deontic logic (von Wright 1968), namely, (for q an arbitrary action):

O_q	q is obligatory
P_q	q is permitted
F_q	q is forbidden.

Using negation, these operators are interdefinable. Permission to do q is equivalent to not being obligated not to do it ($P_q \leftrightarrow \sim O \sim q$), while forbidding q is being obligated not to do q ($F_q \leftrightarrow O \sim q$). Likewise, permission and prohibition (forbidding) are negates ($F_q \leftrightarrow \sim P_q$; $P_q \leftrightarrow \sim F_q$).

To be adequate as a language for bureaucratic procedures, these operators need to include an aspect of contingency (corresponding to the conditions in production rules). Unfortunately, contingency is not straightforward in deontic logic, and a number of proposals appear

(Hilpinen 1981a, 1981b). Note that discretionary actions are those not forbidden, hence permitted. A 'perfect' bureaucracy, in the sense of being completely rationalized and determined, would eliminate permissions entirely. Everything would be either (contingently) obligatory or forbidden.

This is of course a macabre and unworkable design for any human organization. As Norbert Wiener (1967) argued in the early days of computing, such extreme regimentation is an inhuman use of human beings; such activities are not only economically but morally better left to machines.

ORGANIZATIONAL ADAPTATION

Jay Galbraith extends the information processing view of organizations by classifying the environments they face on a two dimensional scale of 'complexity' and 'uncertainty' (Figure 2).

Complexity might be measured in terms of the number of information processing steps (inferences) required to plan the organization's actions.

Uncertainty is essentially the amount of surprise or unpredictability in the environment. This is different from simple contingencies, where the alternatives are foreseen, though the particular outcome is unknown. Uncertainty involves completely surprising events. Thus, as uncertainty increases, planning, even contingent planning, becomes less effective. The organization has to do more and more revision and adaptation while the task is being performed. As an analogy, consider planning a road trip.




Figure 2.

You take along a spare tire, extra oil, etc. for the foreseeable contingencies. Then there is an earthquake, which you didn't expect and so you have to completely revise your plans.

Rationalization, whether by bureaucratic or computer programs, is most effective in situations where complexity is high but uncertainty is low. Surprise requires re-programming, and that tends to be time consuming for either type of software. Left to their own, however, human beings can be quite adaptable. So, Galbraith observes, a counter-strategy in highly uncertain environments is to rely more on individual discretion, rather than trying to pre-program the individual's behavior. This leads to what Burns and Stalker (1961) call 'organic' — as opposed to 'mechanistic' — forms of organization.

This seems to be effective in uncertain environments where complexity is low. However, beyond certain fairly modest levels, unaided human cognition suffers memory limitations and computational biases (Simon 1955, Miller 1956, Tversky and Kahneman 1974). How, then, should an organization which faces an environment that is both complex and uncertain be administered?

It is in response to this question that AI research seems most promising. The appeal is that while bureaucratic procedures are generally written to be deterministic and inflexible, AI problem solving research has led to approaches where numerous heuristics can be tried for a particular problem variant. If one strategy doesn't work, we backtrack and look for another. Several strategies may in fact be satisfactory in which case we can widen the scope to include less consequential problem variables and so provide adaptive, responsive solutions that simple, deterministic bureaucratic methods don't uncover.

MECHANICAL VS MANAGERIAL COGNITION

So far we have considered only the character of the instructions given to the problem processor, automaton vs human administrator. The instructions were expressed in languages L_C and L_B respectively. We now consider the language L_{RW} which these entities process. Typically the input stream includes some description of the problem, while the output stream is a course of action (to be followed by other entities in the organization, whether machine or human or both).

Management texts typically divide the activities of managers into *planning* and *control*. In a planning problem the input is current and predicted information about the external environment and the output is a plan to be followed by subordinate entities (departments, people, machines) in the organization. In a control problem the input is current and predicted facts about the internal environment as compared to an existing plan. The output is a revision to the plan.

This view is quite compatible with the conception of planning in AI. There are however two differences which eventually limit the degree to which AI technology can take over management tasks in the organization. We refer to these as the 'ego' and the 'semantic' problems.

A. The Ego Problem

People have preferences, computers don't. Computers (as we know them) will never prefer chocolate to vanilla. By preference we mean basic or intrinsic values, as opposed to instrumental or intermediate goals. Chess programs, for instance, have intermediate goals leading to the winning of the game. The goal of winning itself, however, is presumed prior to the system design.

The argument here is not absolute, but rather political. We could for instance imagine a robot with high priority heuristics for survival. This might lead down eventually to a sub-goals such as a taste for sweets or a compulsion to win at chess. However, we aren't likely to *allow* such machines to indulge these preferences if they compete with our own. (Note how Asimov's robots (1978) are *programmed* to be socially inferior.) Robot suffrage is not forthcoming.

The converse concept to the social right to have and indulge one's preferences is *responsibility*. The outcome of a computer fraud trial is never to put the computer in jail. Interestingly, not only people but also organizations are granted this social status. A corporation (as well, a sovereign state) has independent legal responsibility; it can sign contracts, can be sued, etc.

The preferences (goals, values) of an organization are generally regarded as deriving from the preferences of individuals. Capitalist economics assumes these to be the values of investors. Socialist economics presumes these are imposed by the society at large. Theories of organization, however, tend to ascribe a larger role to the preferences of people within the organization. Cyert and March (1963) note that the influence of stockholders in large corporations has come to be minimal, and regard the preferences of managers as more significant in a predictive theory. Earlier, bureaucracies were characterized as organizations where the influence of individual preferences was minimized. Managers fill prescribed roles and are substitutable over time. The organization's life is not limited to the life of its members. On the other hand, the mechanistic character of bureaucracy which gives it permanence, also fixes its value structure. Hence railroads, post offices and the military continue to pursue ends that no longer coincide with social interests (Boulding 1978).

In the other extreme, March and Olsen (1979) discuss the nature of organizations where the goals expressed in the organization's formal charter are vague and difficult to measure -- e.g., universities, research institutions, charity organizations, etc. Here the organization's goals are

heavily influenced by those of individual members, and shift in a fluid way in what they call a 'garbage can process'.

Deal and Kennedy (1982) provide an interesting intermediate viewpoint in their concept of 'corporate culture' (see also Peters 1980). In numerous case examples, for instance IBM, General Electric, Dupont, and 'Japan, Inc.', they observe coordinated, cohesive behavior yet without heavy bureaucratic regulation. The differentiating variable, they argue, is that these organizations have built a strong organizational culture which influences and molds individual drives and interests to coincide with the organization at large. Conversely, individual preferences and values also exert influence on those of the organization. The dual membership of the individual in the corporate culture as well as the culture at large ensures that the organization maintains goals and values compatible with its larger social context.

The point is that individual preferences play an important role in the adaptation and goodness-of-fit of the organization to its social environment. While we might conceive of a scenario where a robot or information system also displayed intrinsic preferences, this would be socially inadmissible (and has been in all the science fiction to date). It is of course not the preference itself but the tendency to indulge that preference that matters. Having the right to indulge one's preferences (within socially defined bounds) amounts to political participation, a right still not won by all human beings, let alone robots.

We observed in the beginning of this section that an important function of managers is planning. Planning is also an important AI topic. However, one limitation of AI systems to do organizational planning is in

the selection of the ultimate preferences and values to which the plans are directed. Another limitation, a semantic one, is discussed next.

B. The Semantic Problem

Everyone knows that computational semantics is hard. We argue that for management applications semantics is *impossible*, so long as computers don't have a social life.

Semantics is a rather touchy subject, since there are a number of definitions that circulate and they are rather hard to separate. Generally, semantics is the correspondence between a symbol system (language) and its referents.

In the first section we distinguished between L_C , the language referring to the computer and its operation, from L_{RW} , which referred to the organizational environment. In current terminology this might be phrased as programming language semantics vs database semantics. As before, we attempt to avoid the present debates (e.g., various data management models vs semantic network representations) by skipping over aspects of psychological modeling, retrieval efficiency, etc. and assume that L_{RW} can be characterized as a (first-order) predicate calculus language.

The other advantage of this assumption is that it helps to focus the immense literature on formal semantics without computational distractions. In the predicate calculus (data management and semantic nets as well) we typically make the assumption that semantics follows syntax. That is, the semantics of complex expressions is constructible from the

semantics of its syntactic constituents. (Dowty et al. 1981:Ch. 2). This is Frege's 'Principle of Compositionality'.* The role of the usual logical connectives and quantifiers in constructing the semantics of first order assertions is well studied (van Fraassen 1971). What remains is the semantics of the open vocabulary of the logic, namely individual and predicate names. The approaches at this point divide roughly into two camps, what we will call the *extensional* and *intensional* viewpoints.

Extensional Semantics

The extensional viewpoint is dominant in formal logic, originating mainly from the model theory of Tarski (1956). Here, individual objects are regarded as primitive, leaving generic properties and relationships to be defined set theoretically. An interpretation or *model*, of a given (first order) predicate logic therefore begins with the assumption of a domain of individuals, D , and an interpretation function, F , which maps individual names to individuals in D , 1-place predicates to subsets of D , n -place predicates to relations on D , etc. Hence a model M of a language L has the form

$$M_L = \langle D, F \rangle.$$

This is entirely satisfactory as long as the population of individuals in D can be clearly specified, and they don't change.

* Here we are speaking of formal, constructed languages. The principle of compositionality doesn't always hold in natural language, e.g., for proper nouns like 'Marilyn Monroe' or nominal compounds like 'red herring' where the referent of the expression is not constructable from the referents of its component words.

However, a problem for management applications is that organizations and their environments *do* change. Change is fundamental to economic growth; it can't be ignored. An obvious step is to extend the model to include a time dimension, T, so that D includes all individuals existing at different times. Models of the language are then of the form:

$$M_L = \langle D, T, F \rangle$$

This, however, encounters difficulties when we consider aspects of the *future*. Much of management is concerned with planning. Since there may be a variety of alternate or contingent plans, we must likewise consider multiple futures. This leads to another extension to the model including so-called possible worlds, W, hence adopting models of the form:

$$M_L = \langle D, T, W, F \rangle$$

This is essentially the ontology proposed by Montague (see Dowty et al: 1981, Lee 1981). While this enables a mathematically elegant solution, the question is whether it is still semantics. If semantics is the correspondence between symbols and the world, but if the world is not merely the actual world (past and present) but also future and hypothetical worlds, we have to consider how it is we know about these other worlds.

Strawson (1959) points out that the principle basis for our shared epistemology is reference within a common spatial/temporal framework. Possible worlds are mental constructions, Gedanken experiments. They are outside the framework of external reference and so are questionable as a basis for mutual understanding. We return to this problem shortly.

Intensional Semantics

The intensional viewpoint is more characteristic of the AI paradigm (especially semantic net representations). Here, it is not individual objects that are primitive, but rather generic properties and relationships. Particular objects and events are seen as instances of these generic concepts. For example, we postulate primitive concepts, MALE, FEMALE, SPOUSE, CHILD and from these are able to define the entire vocabulary of kinship relations. Particular cases of family trees, etc. are regarded as 'instantiations' of these generic concepts.

The intensional approach is entirely satisfactory for what we might call idealized or artificial subject domains, where the scope of variation is fixed theoretically or by explicit rules. However, the intensional approach also has difficulties, especially in describing real world domains where no theoretical foundation exists. For example, suppose we want to develop a concept, LEMON. We then seek to elaborate the essential properties of lemons. This might be a property list something like:

COLOR:	YELLOW
SHAPE:	OVAL
TEXTURE:	BUMPY
TASTE:	ACID

The problem, typically, with real world domains is that we can't simply *define* what a LEMON is, but rather our definition has to correspond to what the users of the system conceive lemons to be. Now we run into the so-called 'criterial properties' problem. We want a set of properties that in conjunction uniquely selects out lemons and only lemons from the

various objects in the environment. The problem here is twofold: that too many things qualify (e.g., yellow limes) and the definition excludes atypical lemons (e.g., green lemons, lemons that aren't oval, etc.). Wittgenstein (1953/1958) is a classic elaboration of these difficulties.

There is an interesting relationship between the effectiveness of the intensional approach and the status of the science of that subject domain. Chemistry, for instance, provides a criterial definition for water (as H₂O). Psychology, by contrast, has no criterial definitions for such phenomena as intelligence or creativity.

The problem seems all the worse in the social/economic domains that are most common to management problems. Take for instance the mundane example of chairs. Is there a single physical characteristic that chairs have in common? Consider such examples as rocking chairs, stuffed chairs, bean-bag chairs, plastic inflatable chairs. It seems that what is common to them all is not what they are, but what we do with them, namely sit. But this is no longer an actual property, but rather a propensity or disposition, which leads to similar epistemological difficulties as with possible worlds. (Rescher (1975:Ch.7) comments on dispositional properties and possible worlds.)

A Sociological View of Semantics

Both the extensional and intensional approaches to semantics suffer epistemological difficulties, especially in the social/economic domains typical for management. This leads to an examination of the mechanisms by which we come to know and use the terms of our everyday language.

If we follow the extensional approach, then our main focus will be on our knowledge and identification of individuals (people and objects). This brings attention to the semantics of proper names and the identification codes we assign to machines and other objects. As Kent (1978) points out, these are of fundamental concern in data processing applications, mapping database records to inventory, equipment, personnel, customers, suppliers, etc.

How are these names associated to individuals? In the case of manufactured objects, quite often the identifying name is stamped directly on the object. In the case of names of persons and companies, the identification relies heavily on honest reporting of their names by the entities themselves, e.g., on employment applications, sales orders, etc. The point is that the organization doesn't have to *recognize* these individuals through some collection of identifying properties, it is simply *told*, e.g., "I am John Doe," "Here is the XYZ company."

The point applies much more broadly. Most of what we know about other individuals (people, places, things) that are temporally or geographically distant is what we have been told. The proper name provides a tag to which various characteristics are attached. The names themselves are passed from one person to the next in a series of 'causal chains' of reference, leading back to a direct identification of the individual. Sometimes, in the case of multiple names for the same individual, the causal chains may separate, leading to assertions like

Mark Twain = Samuel Clemens

having an informative content rather than being a tautological identity.

Kripke (1971, 1972) applies this concept of causal chains in a forward fashion in characterizing possible worlds. "Possible worlds are not far-away planets," they are rather *constructed*, based on known, actual references.

Consider, for instance, a scenario beginning with the supposition that Ronald Reagan is bald. The question arises, how do you know it's Ronald Reagan if, in this possible world, he has different properties. (We can exaggerate the case — suppose Ronald Reagan is really a robot, manufactured on Mars, etc. — this is called the problem of 'trans-world identification of individuals'.) Kripke's point is that we don't have to *recognize* Ronald Reagan in this world, we *stipulate* that he is the same in our construction of the scenario. The proper name Ronald Reagan is a 'rigid designator'.

Putnam (1970, 1978) suggests a somewhat similar explanation to our understanding of generic concepts like 'lemon' and 'chair'. Consider the first example of 'lemons'. Being a poor cook, my concept of lemons is fairly rudimentary. I surely couldn't tell a lemon from a yellow lime. Yet I don't often make mistakes in shopping for them. How do I manage? I go to the supermarket and look for the fruit section. There, typically, is a case labeled 'lemons', where I draw my selection. I rely heavily on the supermarket's knowledge to know what lemons are. But how does the supermarket know? They make purchases orders to a distributor requesting shipment of 'lemons'. How does the distributor know? They order 'lemons' from certain fruit growers. How do the fruit growers know? Eventually the chain goes back to a botanist or agronomist who

has certain scientific criteria for lemons.

Now consider the concept, chair. Again we can follow the chain of reference back, this time to certain chair manufacturing companies. But how do they know what a chair is? They *specify* that their products are chairs. Thus one enterprising company may stuff burlap bags with shredded styrofoam and market it as a 'pillow chair'. Another might fold and paint pieces of cardboard selling them as 'throw-away chairs'. The success of their marketing also succeeds in modifying the concept of chair.

The effect of these arguments is to introduce a sociological conception of semantics, what Schwartz (1977) calls the 'new theory of reference'. It gives a convincing account of why semantics is so difficult to do computationally: semantics isn't fuzzy, it's social. For many of our terms, e.g., lemon, chair, the extension of the concept is quite exacting. A thing is a lemon (chair) or it is not. However, the cognition that makes this discrimination is not an individual one, but rather a cooperation of a broad social network. As Putnam observes, we tend to regard words like hand tools that we use individually. For many words, a more fitting metaphor is to compare them to a big ocean liner that requires a crew of hundreds for its operation.

EXPERT SYSTEMS VS DECISION SUPPORT SYSTEMS

Expert systems are typically built to model individual expertise, e.g., a doctor, a travel agent, an automechanic. The view, generally, is of an independently operating problem solver.

Managers don't appear to be experts in this same sense. Mintzberg (1973), in an empirical study of the activities of high level executives, notes that a great portion of managerial activity is spent in communication, observation and data gathering. Moreover, some 70% of their time is spent in informal meetings and committees. Indeed, in this sample, managers only spent about 22% of their time in isolated concentration. The suggestion here is that managers, rather than possessing an individualized expertise, are more like specialized nodes in a larger 'organizational cognition'. Organizations in turn, react and participate in a larger 'social cognition' in their attempts to market new products and/or novel services.

An important part of the manager's activity is to observe and understand changes and trends in the market, the economic, legal and social environments. Much of this is not simply shifts in magnitude on pre-defined dimensional scales. (Were this so, mathematical models would surely have a bigger impact on managerial practice.) Instead, managerial cognition often involves the modification of primitive concepts. For instance, the range of phenomena we call an 'automobile' changes from year to year. Each competitive innovation, each new marketing angle, each special interest group expands and re-organizes the phenomena the manager includes in his/her conceptual framework. And, given that his/her contact with the world is primarily through linguistic interactions, the semantics of organizational language is constantly shifting.

Because mechanical inference relies on a stable, fixed semantics, the utility of an idealized, fully integrated, knowledge-based inference system will be limited to organizations in completely stable environments.

Similar criticisms can be made of bureaucratic rationalization (Lee 1983).

The conclusion to be drawn is that integrated information systems will only be of use for those aspects of the organization's activities where semantic stability can be maintained. This conclusion corresponds to the empirical observations made by Gorry and Scott-Morton (1971), which led to the conception of 'decision support systems' (e.g., Keen and Scott-Morton (1978), Bonczek et al. (1981), Fick and Sprague (1980), Sol (1983)). The underlying idea in the DSS work is to promote the development of technology which, rather than replace human cognition, seeks to assist and augment it. The trend seems to be towards developing DSS 'generators' which provide computational building blocks which can be variously structured for different ad-hoc decision situations.

Interestingly, despite the widely recognized importance of group decision making, nearly all DSS packages are oriented towards assisting the individual manager in isolation. The explanation may be semantic: an individual can *assign* an interpretation to a particular syntactic representation (s)he invents. In a group setting however, the semantics is *negotiated*, and our technology so far seems to have had little effect on these socio-linguistic processes.

SUMMARY REMARKS

The preceding arguments can be summarized in the following statement: we *make* words *mean* what *we* want. Three aspects are emphasized.

We *makes* words mean what we want.

Semantics is plastic. As Tarskian model theory so bluntly points out, the semantics of a language is an interpretation assigned to it. Certain truths (logical truth) are tautologous in that they hold under any interpretation (true in all possible models). In organizational applications, however, we are more concerned with specific interpretations (synthetic truths, true in some models, not true in others). The validity of the inferences drawn depends on the stability of this interpretation. For example,

$$\text{LEMON}(x) \rightarrow \text{YELLOW}(x)$$

is true if in fact all lemons are yellow, but fails if some botanist succeeds in generating a strain with different colors and declares that they, too, are lemons.

We make words mean what we *want*.

Semantic change has a pragmatic component, depending on the interests, preferences and values of its users.

We make words mean what *we* want.

Semantics is plastic, pragmatic, but also the product of social consensus. Indeed, it is not only socially determined, but socially understood.

POST SCRIPT

The purpose of this paper has been mainly to elaborate a problem rather than propose specific solutions. The point certainly has *not* been to discourage further AI research. Rather, it may serve to explain some of the frustration felt in many of attempts at knowledge representation, particularly in managerial applications. As we suggest here, the problem may be overwhelmingly difficult, requiring ultimately a formal explication of all of society. If that is the case, we would do well to seek out more achievable goals and strategies.

Likewise, we have to be careful not to overstate our claims. As pointed out in the beginning, AI is getting market appeal. Big money is shifting. But the people behind those big decisions aren't technicians nor theoreticians. They aren't accustomed to our tendency to extrapolate world shaking implications from toy-sized implementations. They may actually believe us. And the plans for 1984 are in the making now.

ACKNOWLEDGMENTS

The author gratefully acknowledges the stimulating interactions with Steven Kimbrough, Eckehart Köhler and Amilcar Sernadas on these topics. As well, Werner Shimanovich and other members of the 'Vienna Quadrangle' (our humble remake of the Vienna Circle) provided a general background of discussion linking artificial intelligence and formal philosophy.

REFERENCES

- Asimov, I. 1978. *I, Robot*. New York: Fawcett.
- Bonczek, R.H., C.W. Holsapple and A.B. Whinston. 1981. *Foundations of Decision Support Systems*. New York: Academic Press.
- Boulding, K.E. 1978. *Stable Peace*. Austin: University of Texas Press.
- Burns, T., and G.M. Stalker. 1961. *the Management of Innovation*. London: Tavistock.
- Clocksinn, W.F. and C.S. Mellish. 1981. *Programming in Prolog*. New York: Springer-Verlag.
- Coelho, H., J.C. Cotta and L.M. Pereira. 1980. *How to Solve It With Prolog*. 2nd Edition. Lisbon: Laboratório Nacional de Engenharia Civil.
- Cyert, R.M. and J.G. March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs, New Jersey: Prentice-Hall.
- Davis, R., and J. King. 1975. An Overview of Production Systems. Stanford AI Lab Memo AIM-271, Stanford Computer Science Report. STAN-CS-75-524. Stanford, California.
- Deal, T.E. and Kennedy A.A. 1982. *Corporate Cultures*. Reading, Massachusetts: Addison-Wesley.
- Dowty, D.R., R.W. Wall, and S. Peters. 1981. *Introduction to Montague Semantics*. Boston: Reidel.

- Fick, G. and R.H. Sprague, Jr. eds. 1980. *Decision Support Systems: Issues and Challenges*. Oxford: Pergamon Press.
- Galbraith, J. 1973. *Designing Complex Organizations*. Reading, Mass.: Addison-Wesley Publications.
- Galbraith, J. 1977. *Organization Design*. Reading, Massachusetts: Addison-Wesley.
- Gallaire, H., and J. Minker. eds. 1978. *Logic and Data Bases*. New York and London: Plenum Press.
- Gallaire, H., J. Minker and J.M. Nicolas. eds. 1981. *Advances in Data Base Theory* Volume 1. New York and London: Plenum Press.
- Gorry, G.A. and M.S. Scott-Morton. 1971. A Framework for Management Information Systems. *Sloan Management Review*, 13(1):55-70.
- Hilpinen, R. 1981a. ed. *Deontic Logic: Introductory and Systematic Readings*. Dordrecht: D. Reidel.
- Hilpinen, R. 1981b. ed. *New Studies in Deontic Logic*. Dordrecht: D. Reidel.
- Hopcroft, J.E. and J.d. Ullman. 1969. *Formal Languages and their Relation to Automata*. Reading, Massachusetts: Addison-Wesley.
- Jacques, E. 1976. *A General Theory of Bureaucracy*. London: Heinemann.
- Keen, P.G.W. and M.S. Scott-Morton. 1978. *Decision Support Systems*. Reading, Massachusetts: Addison-Wesley.
- Kent, W. 1978. *Data and Reality*. Amsterdam: North-Holland.
- Kowalski, R. 1979a. Algorithm = Logic + Control. *Communications of the ACM*, 22(7):424-436.
- Kowalski, R. 1979b. *Logic for Problem Solving*. New York and Oxford: North Holland.
- Kripke, S. 1971. Identity and Necessity. In M.K. Munitz, ed. *Identity and Individuation.*, pp. 135-164. New York: New York University Press. Also reprinted in: S.P. Schwartz, ed. 1977. *Naming, Necessity, and Natural Kinds*, London: Cornell University Press.
- Kripke, S.A. 1972. Naming and Necessity. In Davidson and Harman, eds., *Semantics of Natural Language* pp.253-355.
- Lee, R.M. 1980. Analyzing Red Tape: The Performative vs Informative Roles of Bureaucratic Documents. WP-80-184. Laxenburg, Austria: International Institute for Applied Systems Analysis.
- Lee, R.M. 1981. CANDID Description of Commercial and Financial Concepts: A Formal Semantics Approach to Knowledge Representation.

- WP-81-162. Laxenburg, Austria: International Institute for Applied Systems Analysis.
- Lee, R.M. 1983. Bureaucracies, Bureaucrats and Information Technology. WP-83-20. Laxenburg, Austria: International Institute for Applied Systems Analysis.
- March, J.G. and J.P. Olsen. 1979. *Ambiguity and Choice in Organizations*. 2. Edition. Bergen: Universitetsforlaget.
- March, J.G. and H.A. Simon. 1958. *Organizations*. New York: Wiley.
- Maslow, A.H. 1943. A Theory of Human Motivation. *Psychological Review*, 50(4):370-96.
- McGregor, D. 1960. *The Human Side of Enterprise*. New York: McGraw-Hill.
- Miller, G.A. 1956. The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *Psychological Review*, 63(2):81-97.
- Mintzberg, H. 1973. *The Nature of Managerial Work*. New York: Harper & Row.
- Peters, T.J. 1980. Management Systems: The Language of Organizational Character and Competence. *Organizational Dynamics*. Summer.
- Putnam, H. 1970. Is Semantics Possible? In: H.E. Keefer and M.K. Munitz, eds, *Language, Belief, and Metaphysics*. New York: State University of New York Press. Also reprinted in: S.P. Schwartz, ed, *Naming, Necessity, and Natural Kinds*, 1977. London: Cornell University.
- Putnam, H. 1978. *Meaning and the Moral Sciences*. Boston: Routledge & Kegan Paul.
- Rescher, N. 1975. *A Theory of Possibility*. Pittsburgh: University of Pittsburgh Press.
- Schwartz, S.P. ed. 1977. *Naming, Necessity, and Natural Kinds*, London: Cornell University Press.
- Simon, H.A. 1955. A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, Vol. LXIX.
- Sol, H.G. 1983. *Processes and Tools for Decision Support*. Proceedings of IFIP/IIASA Working Conference on Processes and Tools for Decision Support, July 19-21, 1982, Laxenburg, Austria. Amsterdam: North-Holland.
- Strawson, P.F. 1959. *Individuals*. Garden City, New York: Anchor Books.
- Tversky, A. and D. Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(September):1124-1131.

- van Fraassen, B.C. 1971. *Formal Semantics and Logic*. New York: Macmillan.
- von Wright, G.H. 1968. An Essay in Deontic Logic and the General Theory of Action. *Acta Philosophica Fennica*, Fasc. XXI, Amsterdam: North-Holland.
- Weber, M.. 1956/1978. *Economy and Society*. Berkeley, California: University of California Press, translated from *Wirtschaft und Gesellschaft*, Tuebingen: J.C.B. Mohr, 1956.
- Wiener, N. 1967. *Human Use of Human Beings: Cybernetics and Society* New York: Avon Books.
- Wittgenstein, L. 1953/1953. *Philosophical Investigations*. Translated by G.E.M. Anscombe. Third Edition. New York: Macmillan.