ON INFORMATION SYSTEM SEMANTICS: EXPERT VS DECISION SUPPORT SYSTEMS

Ronald M. Lee

February 1983
WP-83-19

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
2361 Laxenburg, Austria

CONTENTS

ON INFORMATION SYSTEM SEMANTICS: EXPERT VS DECISION SUPPORT SYSTEMS

Ronald M. Lee*

Abstract: AI will surely have an important impact on the management of future organizations. The problem now is to make some reasonable assessment of its promise and limitations, in order to direct research and development. Here we examine what an 'ideal' knowledge-based management information system could and could not do. The arguments are based on considerations of formal semantics.

Keywords: knowledge representation, information systems, logical databases, logic programming, formal semantics, expert systems, decision support systems

A. INTRODUCTION

The attempt here is to assess the impact of future AI technology on organizational structures (and society). To avoid the weakness of mere opinion, we base the main part of the discussion on formal semantics.

The major influence of computers in organizations at present is in data processing or information systems. These provide access to one or more formatted databases containing data about the organization's plans, commitments and activities. The aspect we concentrate on here is the semantics of the data so maintained.

The importance of semantics is to justify the inferences performed by the system. For example, if we have a general rule

(*)

$$\forall x \quad P(x) \rightarrow Q(x)$$

* International Institute for Applied Systems Analysis, Laxenburg, Austria.

and observe that P(a), then (by modus ponens) we conclude Q(a). However, the truth of (*) depends on the interpretation we assign to it. For example, if P is interpreted as lemon and Q is fruit, the inference holds. However, if Q has the interpretation vegetable, then the assertion is false and the conclusion is incorrect. But these remarks depend on what we mean by 'lemon,' 'fruit,' and 'vegetables.' Thus, information system semantics relies (to a certain extent) on natural language semantics.

For the inferences of the system to be valid, the semantics of its primitive terms must be stable. If Q means 'fruit' one day and 'vegetable' the next, the systems draws inconsistent conclusions.

It is here that we see an eventual limitation applications within AI for organizational information systems. Organizations, to survive economically, must change and adapt. This involves semantic change. To remain useful to the organization, the information system will need to be correspondingly modified. Of interest is whether, using some future technology, these systems will be able to modify themselves, adapting automatically as the organization adapts. The conclusions drawn are pessimistic, as regards complete solutions, considering the semantic foundations of the subject matter.

A more promising direction, we think, is in the perspective of so-called 'decision support systems' (DSS), where the semantic difficulties may be overcome through a cooperative, dialectical relationship between managers and machines.

B. 'KNOWLEDGE-BASED' INFORMATION SYSTEMS

Database management, viewed abstractly and ignoring efficiency considerations, can be regarded as collections of logical assertions, for instance in a first order predicate calculus (see e.g., Gallaire and Minker 1978, Gallaire et al. 1981). Simultaneously, logic programming, based on a Horn clause form of the predicate calculus, is becoming popular as an AI programming language (e.g., Clocksin and Mellish 1981, Coelho et al. 1980, Kowalski 1979).

Extrapolating these trends, we visualize the possibility of a 'knowledge-based' information system. In this idealized image, the programs and data of the entire information system are seen as a giant theorem prover, providing logically derived conclusions from observed facts about the organization and its environment.

The question, then, is what are the possibilities of such a system? Would management cognition be superceded?

C. FORMAL SEMANTICS

An advantage of characterizing an information system as a theorem prover operating on predicate calculus assertions, is that it brings the substantial literature on formal semantics to bear.

In analyzing the semantics of a formal (predicate calculus) language, the assumption is generally made 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 predicate names and logical constants. The approaches at this point divide roughly into two camps, what we will call the *extensional* and *intensional* viewpoints.

1. 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*, <D,F>, 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 logical constants to individuals in D, 1-place predicates to subsets of D, n-place predicates to relations on D, etc.

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

$$M_L = <D, T, F>$$

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:

$$M_L = <D, T, W, F>$$

* 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 referrent of the expression is not constructable from the referents of it's component words.

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.

2. Intensional Semantics

The intensional viewpoint is I think 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 is to be found.

For example, suppose we want to develop a concept, LEMON. Now we 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 H_2O). 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. Consider 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.)

D. 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 Clemons"

having an informative content rather than a tautological identity.

Kripke (1971, 1972) applies this concept of causal chains in a forward fashion in characterizing possible worlds. "Possible worlds are not faraway 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 descrimination 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.

E. 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 predefined dimensional scales. (Were this so, mathematical models would surely have a bigger impact on managerial practice.)

Instead, it 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 1980).

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 basic idea here 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.

F. CONCLUSION

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.

REFERENCES

Bonczek, R.H., C.W. Holsapple and A.B. Whinston. 1981. *Foundations of Decision Support Systems*. New York: Academic Press.

Clocksin, 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.

Dowty, D.R., R.E. Wall and S. Peters. 1981. *Introduction to Montague Semantics*. Boston: D. Reidel Publishing Company.

Fick, G. and R.H. Sprague, Jr. eds. 1980. *Decision Support Systems: Issues and Challenges*. Oxford: Pergamon Press.

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.

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. 1979. *Logic for Problem Solving*. New York and Oxford: North Holland.

Kripke, S.A. 1972. Naming and Necessity. In Davidson and Harman, eds., *Semantics of Natural Language* pp.253-355. Dordrecht: D. Reidel.

Kripke, S. 1971. Identity and Necessity. In M.K. Munitz, ed. *Identity and Individuation*., pp. 135-164. New York: New York University Press. Reprinted in: Schwartz (1977).

Lee, R.M. 1980. Bureaucracies, Bureaucrats and Technology. WP-80-186. 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.

Mintzberg, H. 1973. *The Nature of Managerial Work*. New York: Harper & Row.

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. Reprinted in Schwartz (1977).

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.

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.

Tarski, A. 1956. The Concept of Truth in Formalized Languages. In A. Tarski, *Logic, Semantics, Metamathematics* (translated by J.H. Woodger. Originally presented to the Warsaw Scientific Society, March 1931 in Polish. Oxford: Clarendon Press.

van Fraassen, B.C. 1971. *Formal Semantics and Logic*. New York: Macmillan.

Wittgenstein, L. 1953/1958. *Philosophical Investigations*, translated by G.E.M. Anscombe. Third Edition. New York: Macmillan.