The corporation as mind: Lessons for AI

Stephen Regoczei
Computer Studies, Trent University
Peterborough, Ontario, CANADA K9J 7B8

Graeme Hirst
Computer Science, University of Toronto
Toronto, Ontario, CANADA M5S 1A4
gh@ai.toronto.edu

The mental life of corporations

The most successful example of distributed artificial intelligence ever constructed is the modern corporation. It is our view that a corporation is an instance of AI—a generalized, distributed AI system (GDAIS). We believe that this more general view carries lessons for the construction of computer-based intelligence.

We use the term *corporation* to include not only large companies, but also similar organizations such as government departments. Such organizations are typically composed of small management groups that, in turn, are organized into larger sections or divisions. The result is usually a tree structure, or something close to that.

The hardware from which this structure is built is of two kinds: equipment and people. The equipment includes all the corporation's buildings, machinery, furniture, and so on. The people are its employees. But, we submit, the intelligence in the corporation lies not so much in the people as in the other components of its hierarchy. Each management group, each corporate division, and, indeed, the corporation itself may be regarded as an abstract intelligent agent: an entity with its own 'mind', a mind that the other components influence through symbolic communication.

At the bottom level of the mental structure of the corporation is the minds of the people it employs. But, even though these people are pretty smart (on a cosmic scale, including rocks, computers, bacteria, and lizards, an IQ of even 100 is remarkably high), usually only a small portion of their intelligence gets used in forming the corporate mind.

Why is this so? Because, strictly speaking, it is not true that a corporation employs a *person*. Rather, what it employs is a *role* of that person. Now, people have, for the most part, general-purpose skills and intelligence. But when they are employed in a particular corporate role, it is usually only for a small fraction of their abilities—especially in the more-numerous lower-level roles. An assembly-line worker might be able to design the cars just as well (or better) than those who do, but since she is not employed to do that, she doesn't.

A corporation is thus an intelligent entity formed from relatively unintelligent components.[1] So on this view, it can be seen as an analogue of Minsky's (1986) 'society of mind' model of intelligence.

This analogy cannot be taken too far, however. In comparison to a corporation, Minsky's model is much too impoverished and can never, we submit, be intelligent the way a corporation is. There are (at least) four important properties of a corporation that are missing from the 'society of mind':

1. A corporation has a structure; it is not just a random collection of agents.
2. A corporation has a culture and a 'tradition'; no component of the corporation, nor the corporation itself, is without a history.
3. The agents in a corporation are situated. They have a knowledge of their local environment, both symbolic and physical, within the corporate structure and culture.
4. The agents in a corporation engage in complex symbolic communication; they don't merely send simple signals to one another.

For while a corporation is an *artificial* intelligence, it is a *natural* society—unlike Minsky's model. A natural society is not just a set of people; it is a set of people with a culture (and this includes corporations, as it has become fashionable to observe). A culture, here, can be thought of as a knowledge base of facts, predispositions of interest, and ways of doing things. It acts, in effect, as a generalized abstract agent that communicates with both individual and group agents in the society to give them a mental 'map' of the society within which they are situated. But Minsky's society of mind has no such thing. Our suggestion is that, as we start to design artificial societies of artificially intelligent agents, the question of social culture does indeed arise. Distributed AI needs *an ethnography of constructible societies.*

Intelligence in a group agent thus requires a very rich kind of inter-agent communication. Our view of situated, symbol-using, 'cogniting' agents remedies some of the shortcomings of inflexible, non-adaptive models. No realistic model of social interaction can be built in DAI without taking symbolic communication into account.

[1] This might be seen as an overly cynical view of corporations, but it is, we submit, not unrealistic. Many of the current ideas for making corporations better (*i.e.,* more competitive or whatever) revolve around making better use of their people—in effect, boosting the corporate intelligence by using more employee intelligence. This is a tacit recognition that employees' intelligence is not normally used to its full extent. In addition, our view explains how perfectly reasonable people can turn into the stereotypical dumb, unhelpful clerks, assistants, and bureaucrats that we all encounter from time to time in corporate structures that are particularly unrewarding of, or even intolerant of, the exercise of excess intelligence by their employees. It also explains how committees of intelligent people come to make stupid decisions.

Works that have influenced us

We have been mostly influenced by developments in General Semantics and General Systems Theory, and constructivism / perspectivism / conceptualism (taken seriously).

- P.L. Berger and T. Luckmann (1967). *The social construction of reality: A treatise in the sociology of knowledge.* Doubleday Anchor.

The title is unneccessarily provocative. It should be *The construction of what counts as reality in a social setting.* It is a powerful document, demonstrating the constructivist and perspectivist positions. Positivists and 'hard science' types find it painful reading.

- Gerald Holton (1988). *The thematic origins of scientific thought: Kepler to Einstein* (Revised edition). Harvard University Press.

There is more to scientific work than a dispassionate, objective search for truth. This book is gilt-edged evidence from a field as 'unflaky' and respectable as physics that the personal thematic preferences of scientists have a decisive effect on the theories they forge. We need a similar work for the field of AI, and computer science in general.

- Anatol Rapoport (1953). *Operational philosophy: Integrating knowledge and action.* Harper.

Unites the concerns of operationalism with the insights of General Semantics. A seminal work, and a superb introduction to some of the most crucial problems of knowledge, discourse, and action in the social context.

- S.I. Hayakawa (1941). *Language in action: A guide to accurate thinking.* Harcourt Brace.

In the tradition of General Semantics, Hayakawa demonstrates how the creation of different domains of discourse and the attachment of different concept clusters to the same piece of text not only alters meaning but leads to different courses of action.

- Bill Kent (1979). *Data and reality.* North-Holland.

The map is not the territory. Kent struggles with the paradoxes of database design as a modeling activity. Highlights how data-processing models fall far short of the reality they claim to be modeling. The issues he raises are still current.

- Mary Douglas (1986). *How institutions think.* Syracuse University Press.

The title says it all. The key issue is this: In what sense can we operationalize the vague notion of the group mind? (*cf* Rapoport for operationalizability and Connerton for a similar investigation.)

- Benedict Anderson (1983). *Imagined communities: Reflections on the origin and spread of nationalism.* Verso.

Nationalism, ethnicity, political identity, and ideological stance are typical social constructs that do not seem 'real'—hence the term *imagined*—and yet are very powerful. Knowledge and action in social and organizational contexts focuses on the creation and use of such constructs.

- Paul Connerton (1989). *How societies remember.* Cambridge University Press.

Memory, record keeping, and construct creation and transmission as social activities (*cf* Douglas).

- Ludwig von Bertalanffy (1969). *General systems theory.*

The classic in the field. Badly in need of updating. Recent work is centering upon systems with large software components. We have come a long way, but it is good to know the genesis.

- Martin Gardner (1962). *Fads and fallacies in the name of science.* Dover.

A wonderful collection of case studies on knowledge and action in social and organizational contexts.

- John Sowa (1984). *Conceptual structures: Information processing in mind and machine.* Addison-Wesley.

A significant work for both the conceptualist and the perspectivist approach. The conceptual-graph notation is most helpful for discourse analysis. Needs a supplementary work to explore the constructivist implications, and to address the operationalizability of the conceptual analysis techniqes that permeate the book.

- Marvin Minsky (1986). *The society of mind.* New York: Simon & Schuster.

No comment.

Some of our results so far

Hirst, Graeme (1987). *Semantic interpretation and the resolution of ambiguity.* Cambridge University Press.

Hirst, Graeme (1991). "Existence assumptions in knowledge representation." *Artificial intelligence*, **49**, May 1991, 199–242.

Regoczei, Stephen and Hirst, Graeme (1989). "On 'Extracting knowledge from text': Modelling the architecture of language users." *Proceedings, Third European Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Paris, July 1989, 196–211. [Also technical report CSRI-225, Computer Systems Research Institute, University of Toronto.]

Regoczei, Stephen and Hirst, Graeme (1990). "The meaning triangle as a tool for the acquisition of abstract, conceptual knowledge." *International journal of man–machine studies*, **33**(5), 505–520.

Regoczei, Stephen and Plantinga, Edwin (1987). "Creating the domain of discourse: Ontology and inventory." *International journal of man–machine studies*, **27**(3), 235–250. Reprinted in: Boose, John H. and Gaines, Brian R. (editors). *Knowledge acquisition tools for expert systems.* London: Academic Press, 1988, 293–308.

Interesting questions

1. The feasibility of redrawing disciplinary boundaries

Is it time to cut the umbilical chord between artificial intelligence and the computer field? Can we refocus AI research to include other systems with a large cognitive component that are intelligent yet artificial, such as cultures, societies, institutions, and corporations?

Can we take Herbert Simon's 'sciences of the artificial' seriously? Can a study of general distributed AI systems (GDAISs) be formulated as one of the sciences of the artificial? In this world of constructs, dominated by administrators, lawyers, and accountants, can a study of constructs provide us with insight on how to design better general distributed AI systems?

Can we develop better techniques to use perspectives and new conceptual constructs for designing better GDAISs? Can the development of public policy and the analysis of social systems be improved by using insights from AI systems research?

2. The feasibility of agent and meaning modeling

Can we build operationalizable, computational models of roles, individuals, and abstract corporate agents as Durkheimian social units? Can we use extensions and generalizations of the meaning triangle to model situated, symbol-using, cogniting agents?

An operationalizable model might still be merely theoretical, and so the question arises: can we *implement* operationalizable models of situated cogniting agents?

Can we extend the techniques of object-oriented programming to implement models of interacting, situated, cogniting, symbol-using agents? Is such an agent-oriented design and programming paradigm feasible? Given that we presently conceive of the architectures of objects as abstract data types, what kinds of architectures would be *desirable* for such agents?

3. Plato's heaven versus experiential actualities

Is it correct to say that there is no knowledge without a knower? Is it true that there is no discourse without agents who participate in the discourse?

Is there such a thing as the true meaning of a sentence or sign or symbol or proposition? Is it true that meaning is socially negotiated? Are *all* consensual domains of discourse created through the harmonization of mental models?

Are the power-based discourses described by Michel Foucault social constructs? Can we shape, alter, or redesign such discourses by altering the software that 'runs' on the social unit? If Foucault's discourses are power-based, what sort of power base do we need to effect this kind of DiscourseHacking [sic]?

Are concepts and mathematical ideas agent-created constructs? Or are we prisoners of Plato's cave, only barely being able to intuit a few pale images of the pure forms of Plato's heaven? Was Plato merely struggling with the phenomenon of pattern recognition or category formation when he introduced his notion of pure form? Are Plato's forms nothing more than idealizations of the content of a group mind? Is there a need to consider a Platonic heaven when considerations of a group mind are sufficient? Is this a good place to apply Occam's razor? If so, where should we make the cut? Is the group mind already too metaphysical for us to take seriously? Do institutions think? Do societies remember? Or are such expressions merely category mistakes or hyperextended metaphors?

4. Discourse structure and natural language understanding at social and organizational levels[2]

What should a discourse model model at the social and organization level? How do institutions, corporations, ethnic groups, and nations talk to each other? What does it mean to 'change the mind' of an ethnic group through means such as propaganda, advertising, or the managing of news coverage?

What are the components of a comprehensive discourse model in the social and organizational context? Evidently the mental models of the group agents, and the conceptual structures they use, would override any interpretation of mere text, or of visual, auditory, or tactile symbols, would they not? Is the notion of concept cluster attachment (Regoczei and Hirst 1989) rich enough to capture the arbitrariness of interpretation during the communication process?

How should social discourse models be represented? If the modeling of agents, their cognitive architectures, their cognitive content, and their cognitive evolution are all essential for the description of discourse in the social and organizational context, can we ever go back to simple story grammars, and other low-level descriptive tools?

5. Is the grand vision feasible?

Is it possible to construct our theories within the following conceptual framework: An organization (government agency, ethnic group, political unit, institution, corporation, business) can be looked upon as a generalized distributed AI system with a hierarchical structure.

Can we look upon the higher nodes of the tree (departments, divisions, subsidiaries) as abstract agents in their own right? Can we consider the leaves of the tree (employees, members of a political party) not as biological individuals but as roles? Can we model these roles as abstract agents?

Are social units decomposable into hardware and software components? Can we reprogram a social unit by changing the software that 'runs' on it?

[2] It is worth noting that the questions posed in the call for papers of the Discourse Structure Symposium in this series could be re-asked in the social and organizational context. Current work in discourse usually assumes individual agents rather than group agents such as corporations or nations, but the questions still make sense when this assumption is dropped.