Reprogramming the Software That Runs on Organizations: Social Agents as Knowledge-Based Computational Automata

Stephen B. Regoczei

Computer Studies
Trent University
Peterborough, Ontario
Canada K9J 7B8
Phone: (705)748-1495
FAX: (705)748-1625
E-mail: stephen@ai.toronto.edu

Draft
Version 2.0

(Preliminary version of a paper submitted to the AAAI-93 Workshop: AI and Theories of Groups and Organizations: Conceptual and Empirical Research, Washington, DC, 11-16 July 1993.)

business law, was a bit more verbose. He mused out loud:

> *"Organizations? I don't know about organizations. "Organization" is just a wishy-washy term for a corporation. I do know about corporations. Corporations, as the term suggests, are bodies. They are legally created persons—artificial persons, in fact. And they are brought into being through legal documents."*

We should add, that our intrepid explorers did not encounter the ghost of Durkheim. Had they done so, he would have explained to them, at great length, that an organization is really and abstract social agent. George Herbert Mead also did not appear, but we know from other evidence that he would have opted for a role-based explanation. In other words, it does not matter what an organization is; the important consideration is the roles that the organization plays.

What may be the moral of this story? What is an organization?

3. WHAT IS AN ORGANIZATION?

Gareth Morgan, in his seminal book "Images of Organization", opts to describe organizations from different perspectives. It is a challenge to read through this list and to try to decide how one might be able to model each of the versions of the organization using computational techniques. Morgan's list, somewhat paraphrased, is the following:

1. Organizations are mechanical machines.

2. Organizations are situated organisms, subject to Darwinian processes.

3. Organizations are cybernetic, learning, brain-like structures, engaged in a process of self-organization.

4. Organizations are cultures that create their own social reality.

5. Organizations are political systems, laden with conflicting interests and struggling power structures.

6. Organizations are psychic prisons, where the psychodrama of subconscious processes and conflicting ideologies unfolds.

7. Organizations are manifestations of dynamic change, flux, or transformation—in other words, self-producing autopoietic systems (not to be confused with self-organization above).

8. Organizations are instruments of domination and ugly-faced exploitation.

To this list, we may add four more perspectives mentioned above:

9. Organizations are computers.

10. Organizations are legally-created artificial persons.

11. Organizations are abstract social agents. (Durkheim)

12. Organizations are the bundles of roles they play. (Mead)

There are several problems with Morgan's list. One is that many of his perspectives interact in non-trivial ways. For example, organisms and manifestations of flux and change both are to be self-organizing in some sense. But how? Also, many of Morgan's categories are subcategories of systems in general. Yet some, *e.g.* version 3, are far too heterogeneous and complex from a computational point of view. Besides, is Morgan's list, even expanded by the views of Durkheim *et al.*, really as comprehensive as it seems?

4. WHAT CAN WE MODEL WITH KNOWLEDGE-BASED COMPUTATIONAL AUTOMATA?

In earlier papers we argued for a view of the corporation as a knowledge-based, generalized, distributed AI system (GDAIS). (Regoczei & Hirst, Regoczei) We construed the corporation as an "intelligent" entity formed from relatively unintelligent components, namely employees playing restricted job function roles. Yet a corporation can not be modelled as an analogue of Minsky's (1986) "society of mind" because:

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 do not merely send simple signals to each other.

Thus a corporation is laden with cultural software, and it is this software that has to be changed in other to, for example, make the corporation into a better functioning unit.

The significant role of the cultural component prompted us to propose an "ethnography of constructible societies". This idea is complementary to S.L. Star's interesting suggestion of replacing the Turing test with the "Durkheim test". Our ethnographic suggestion is not a rival one. We really do not believe that there should be a sharp division between ethnography and sociology.

5. COMPUTERS, AND HENCE ORGANIZATIONS, AS SLAVES

While our main concern is to go from organizational theory to computational paradigm, let us reverse this direction. Let us look at typical computers today and see what type of an organization they may correspond to. The result, I think is surprising. We are so used to ordering the computer around, issuing it commands, totally disregarding its feelings, if any, that we do not notice how unlike our treatment of the computer is compared to the treatment we should give to other people, or even the way many of us interact with our own automobiles.

The conclusion is inescapable: there is no current organizational type that corresponds to the current computer, although there used to be slave-based organizations in the past.

Tongue in cheek, we could say that this revelation almost cries out for a computer liberation movement. In fact, reading Winograd and Flores's book for the first time, I was struck by the paradigm change advocated therein. It seemed to me that the hidden agenda of the authors was to liberate computers from the indignity of being ordered around. Issuing commands, whether to people or computers, was authoritarian, paternalistic, domineering and hierarchy-based. The computer, rather, was to be encouraged to "do it's own thing". My more bilious East Coast friends were quite ready to characterise the whole suggestion as typical California thinking.

Translating these issues into the discourse of organizational theory, we may add to our list:

13. Organizations are the willing, obedient slaves of the manager/owner.

and the contrasting view:

14. Organizations are the wise cooperative helpmates of the manager/owner.

To elaborate on the distinction between organization-as-slave versus organization-as-cooperative helpmate, we could picture the manager as the driver of a car. If the car is an ordinary standard model, guiding it requires tight, hands-on control. On the other hand, if our fantasy car is more like KITT from Knight Rider, we can rely upon the vehicle to be not only cooperative, but at times wiser than the manager.

And while we are expanding our list, we probably should also include the "none-of-the-above" category:

15. Organizations are cooperatives—everybody is a manager, owner, and helper at the same time working for the benefit of all the publics and stakeholders of the corporation.

This is an ideal that would scarcely show up in an organizational theory textbook such as Morgan's. It is this comment that justifies the position of the business law-type above, that "organization" is a euphemism for the more dirty-sounding "corporation".

But even this list above does not exhaust the obvious candidates. Drucker's organization as a task/goal oriented entity, or Charles Handy's suggestions of the Shamrock, the Federal, and the Triple-I organizations immediately come to mind. And yet there are still other types of organizations. In particular, franchise operations seem to be different.

16. Local franchise organizations are software clones of the franchiser.

Of all the metaphors, perhaps this is the one that has the most immediate computational interpretation. What does one purchase when one purchases a franchise? A software package that comes with the right of being used in a certain area. The software metaphor is directly applicable. The franchiser develops the software and sells it to run on the hardware locally. The local hardware consists of physical plant, and "meatware" that can be hired at minium wages and trained. The software for training is available, and a lot of the software is embodied in the constraints and affordances of the equipment. In other words, the software is Gibsonion, embodied software.

6. IS THE ORGANIZATION A TURING MACHINE?

Can organizations be construed as Turing Machines? Obviously not. The Turing Machine metaphor simply does not fit organizations. Are von Neumann machines better suited? Yes, if we adopt the older notions/images of the organization: the manager issues orders and the corporation executes the commands.

But a better metaphor would be found in object-oriented software or systems. Let us call these systems, for the sake of terminological uniformity, Object machines. This is fine, as far as it goes, but the components of an organization such as roles, people, departments, and subdivisions are more like cogniting agents themselves than dumb molecules bouncing around in a box. If we take this note of caution seriously, we probably would want to make the organization into an Agent machine—whatever its architecture and implementation may be like.

Given that we want Agent machines that are situated, and act and interact in a local environment, do we have obvious concrete examples? Yes. For example, anthills, as described by Edward O. Wilson, sound like large, distributed computational devices that run using biochemical mechanisms such as pheromones. I suppose we could call such devices Wilson machines. Having modelled their architecture, we could contrast Wilson machines and Turing machines as computational devices.

Contrasting a Wilson machine and a Turing machine, even in a casual, informal and intuitive way, is quite an eye opener, but there is no need to stop here. Minsky's "Society of Mind" can be construed as a Minsky machine, whereas Herbert Simon's programmable organization, I suppose, would be a Simon machine. Malinowski machines would have an ethnographic flavour. Schutz machines would worry about the "life world". Durkheim machines would consist of hierarchies of nested, interacting social agents. It makes sense. If we expect the Turing machine to pass the Turing test, then we would need a Durkheim machine to pass Star's "Durkheim test". If there is competition for survival in our society, then the participating agents are, I suppose, Darwin machines, and so on.

7. PROGRAMMING ORGANIZATIONS

The notion of programming organizations in the way we would program a von Neumann computer is nothing new. March & Simon, as early as 1958, in their book on organizations, explicitly broached the programming and reprogramming of administrative decision-making systems.

Could the March & Simon idea of organizations as programmable automata—but perhaps not yet cogniting or knowledge-laden systems—be exploited? With the clarity of hindsight, it seems clear that it could not have been developed at the time when it was proposed. What was missing? On the organizational theory side, there was no emphasis or appreciation of the importance of corporate culture. On the computational side, knowledge-based systems (Feigenbaum) have not yet become a dominant paradigm in artificial intelligence work.

Nevertheless, the hardware-software decomposability implicit in the March & Simon idea can serve as a powerful stimulus for our current thinking regarding the reprogrammability of organizations. In particular, the so-called "learning organization" could be viewed in terms of the learning machine metaphor.

8. WHAT ARE THE MAIN ISSUES?

The main issue is to try not to fish on the road, or to look for the contact lens under the lamppost where the light is better. It seems necessary to take a more perspectivist approach to the reading, understanding, and computational modelling of organizations. As Morgan states:

> *Any realistic approach to organizational analysis must start from the premise that organizations can be many things at one and the same time. A machinelike organization designed to achieve specific goals can simultaneously be: a species of [organism] that is able to survive in certain environments but not [in] others; an information-processing system that is skilled in certain kinds of learning but not in others; a cultural milieu characterized by distinctive values, beliefs, and social practices; a political system where people jostle to further their own ends; an arena where various subconscious or ideological struggles take place; an artifact or manifestation of a deeper process of social change; an instrument used by one group of people to exploit and dominate others; and so on. Though managers and organization theorists often attempt to override this complexity by assuming that organizations are ultimately rational phenomena that must be understood with reference to their goals or objectives, this assumption often gets in the way of realistic analysis. If one truly wishes to understand an organization, it is much wiser to start from the premise that organizations are complex, ambiguous, and paradoxical.* (Morgan, p.321ff)

The inclusion of ambiguity, nuance, inconsistency and downright contradiction seems incompatible with the ideals of a purely rational computational paradigm. But already in the case of knowledge acquisition for expert systems we had to confront the fact that experts are not as pristine in their thought habits as they appear to be. The manipulation of inconsistent, or seemingly inconsistent knowledge without too much effort is one of the hallmarks of an expert, not to speak of lies and obfuscation for political purposes. We take a strangely protective and paternalistic approach to our computer systems when we refrain from telling them how it is really done at times: ambiguity, leaps of faith, solid gut feel, and obfuscation is, as a matter of fact, part of the tool kit. Our moralistic streak, as individuals, is built into our computational models.

Nuance is a different and more subtle issue. The old AI paradigm of problem solving as search within a predefined solution space gave early successes but eventually revealed its fundamental limitations. The burden placed on the designer of the system was too great. The designer's job was to define the solutions space. Often this could not be done realistically; the resulting microworlds seemed like a parody of what was really going on. (Winograd & Flores) The ability to detect subtle differences—subtle nuances—and to use these differences to creatively form new concepts and new scripts (Schank) seems to be a basic feature of knowledge-based, adaptive systems capable of deutero-learning (Casti). (Include material on nuance from Briggs & Peat.) The classification tasks are also more difficult than we used to imagine (Lakoff).

Therefore, Morgan's statement that *"organizations are complex, ambiguous, and paradoxical"* should not be interpreted in a pessimistic way. It does not mean that realistic computational models of organizations are not possible. It may only mean that an approach in which the von Neumann architecture, first order logic, and physical symbol systems are the dominant landmarks is not the best route to take. I personally would like to see not only a system that plays go, but also a system that "solves" koans. After all, our ability to cope with contradictions and paradoxes is very much part of our cognitive furniture. Let us model things as they are. Boulding makes this plea in The Image—let's have the right image! Let us give up producing nothing but lobotomized computer systems. Work on systems that "get emotional" and engage in other seemingly "irrational" but quite effective (in the Darwinian sense) behaviour give indication of one possible avenue to pursue. "Getting emotional" provides special advantages because "reasoning" is very resource-intensive, and in resource-constrained, real-time situations robots may be better off just to "get angry" or "experience fear", and to base their actions and responses on such subconceptual mental states. (Clark Eliot) This work could be looked upon as attempting to provide partial answers to the Heideggerian "thrownness", "at-handedness", "horizon", and "breakdown" issues that Winograd and Flores raise. And if all of this seems too far removed from traditional AI concerns, there are indications that something as basic as natural language understanding may only be feasible through resource-constrained, "sloppy", and satisficing systems that are much more like Rodney Brooks's ants than Newell and Simon's physical symbol systems. (Corriveau)

Morgan continues on a hopeful note, and I think that workers in computational modelling also have grounds for optimism.

> *Fortunately, the kind of metaphorical analysis developed [here] provides us with an effective means of dealing with this complexity. For it shows us how we can open our thought processes so that we can read the same situation from multiple perspectives in a critical and informed way.* (Morgan, p.322)

To reinterpret this point, the computational model is not there to replace, but to enhance the

human person. The computational model is there to enrich our thought processes and to help us to be creative. It should help us to read the same situation from multiple perspectives in a critical and better informed way. Stefik emphasises a similar point in his paper "The Next Knowledge Medium".

> *The most widely understood goal of artificial intelligence is to understand and build autonomous, intelligent, thinking machines. A perhaps larger opportunity and complementary goal is to understand and build an interactive knowledge medium.* (Stefik, p.315)

Those who despair at seeing the limited potential of AI may be doing nothing more that focusing on the wrong problem. The task is not to replace human cognition with cogniting robots. The task is to enhance human capabilities. Stefik's prediction is that this will come about not by building AI systems in the traditional sense but by creating new, interactive knowledge mediums that will help people function better as people. (Stefik, Regoczei) My prediction here is that, as Senge emphasises, computational models are there to help people function better within organizations, not to replace human decision making and administrative judgment. (Senge) Richer computational models of organizations, following Morgan's metaphors (as if they were requirements specifications for future systems to be built) should enhance our ability to think about organizations, to do "what-if" analysis more realistically, and to enhance our ability to redesign and "reprogram" organizations as necessary.

The larger social issues are inescapable. Corporations seem to be the dominant life-form on Earth today. They are truly Darwin machines. Simulation models indicate that the limits-to-growth carrying capacity is being strained. (Meadows *et al.*) If we, as human beings, do not have better tools to "reprogram" corporations, organizations, and impersonal socio-cultural systems, then the future does not look promising. Broadening our scope within AI by providing better computational models of organizations may turn out to play a pivotal role in the evolution of humanity.

9. CONCLUSION

Examining computational models of organizational theories is an exercise in cross-cultural mediation. To thrive in the organizational domain of discourse, we need both worlds. It is not an either/or decision: we need to consider both cogniting agents and computational automata. We need both physical symbol systems and physical grounded systems. We need Rodney Brooks's ants with culture, knowledge, and a social conscience. Organizations whose software is not functioning as it should, may have to be reprogrammed.

10. ACKNOWLEDGMENTS

I am grateful to Graeme Hirst for his continuing encouragement and helpful comments. Preparation of this paper was supported by a generous grant from The Canadian Newt and Eft Research Fund for Knowledge Acquisition Studies.

Bibliography: *(abbreviated version)*

Casti, J.L.: Review of Rosen, 1991, *Bulletin of the Santa Fe Institute*, v.7, 2, 1992.

Corriveau, Jean-Pierre: Time-Constrained Memory for Reader-Based Text Comprehension, Ph.D. Thesis, University of Toronto, 1991.

Gasser, L. & Huhns, M.N. (eds.): Distributed Artificial Intelligence, V.II, 1989.

Harris, Marvin: Cultural Materialism, 1979.

Huberman, B.A. (ed.): The Ecology of Computation, 1988.

Morgan, Gareth: Images of Organization, 1986.

Regoczei, S.: "Agent and meaning modelling as a comprehensive general interpretation of knowledge acquisition", AAAI Spring Symposium, 1992.

Regoczei, S. & G. Hirst: "The corporation as mind: Lessons for AI", AAAI Fall Symposium, 1991.

Rosen, R.: Life Itself, 1991.

Schank, R.: The Creative Attitude, 1988.

Senge, Peter: The Fifth Discipline, 1992.

Star, S.L.: "The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving" in Gasser&Huhns, 1989.

Stefik, Mark: "The next knowledge medium" in Huberman, 1988.