Semiotic Agent Models for Simulating Socio-Technical Organizations

Cliff Joslyn *

Prepared for the DS Project, PSL/NMSU
September, 1999

Contents

*Distributed Knowledge Systems Team, Computer Research and Applications Group (CIC-3), Mail Stop B265, Los Alamos National Laboratory, Los Alamos, NM 87545, USA, `joslyn@lanl.gov`, `www.c3.lanl.gov/~joslyn`, (505) 667-9096.

Abstract

In this document we survey the issues surrounding the approach to modeling Socio-Technical Organizations (STOs) based on *semiotic agents*. We begin with a brief survey of agent concepts and applicability, and then introduce the semiotic approach specifically. We then introduce and describe STOs. While our approach is general to STOs, we are especially interested in exploring particular kinds of STOs such as 911/ERS, search and resuce operations, military organizations, and others which are characterized by a hierarchical or Structured Command Organizations (SCOs). We conclude by introducing SCOs and considering some of their special considerations.

1 Introduction

The world around is becoming composed of systems involving computer-human interaction of unprecedented scale and complexity. The modern environment is an interlocking collective of large numbers of groups of people interacting with computer systems, and which themselves interact with a variety of physical systems to maintain them under conditions of good control. The vast complexity and quantity of information involved makes simulation approaches necessary, and yet the existing formalisms available for simulation are not sufficient to reflect their full characteristics. In particular, simulations built on strict formalisms such as discrete-event systems or hybrid control cannot capture the inherent freedom available to humans interacting with such systems; and simulations built on classical rule-based Artificial Intelligence (AI) approaches are too brittle and specific to allow for the emergent phenomena which characterize such systems.

In this document we argue that an agent modeling approach between collective automata systems such as used in Artificial Life (ALife) and full AI may provide a robust capability to simulate human-machine interaction at the collective level. We call this approach **semiotic**, as it focuses on the use and communication of symbols by and between agents and their environments.

Below we first discuss the current state of the use of agents with respect to both decision theory and modeling and simulation. We then outline the semiotic approach to agent simulation, and then our view of the Socio-Technical Organizations (STOs) intended to be modeled.

While our approach is general to STOs, we are especially interested in exploring particular kinds of STOs such as 911/ERS, search and resuce operations, military organizations, and others which are characterized by a hierarchical structure of command. We characterize these as Structured Command Organizations (SCOs), and conclude by considering some of their special characteristics.

We note that this document is written in close conjunction with Rocha's "Review of Agent Models: Encounters, Strategies, Learning, and Evolution", also provided this distribution. We will avoid duplicate development of material included there, but rather cite it frequently, referring to is as [Rocha p. x].

2 Agents for Simulation

Recent developments in software engineering, artificial intelligence, complex systems, and simulation science have placed an increasing emphasis on concepts of autonomous and/or intelligent **agents** as the hallmark of a new movement in information systems. The history of computer science has seen a "march of paradigms", as programming theory has moved from procedural through functional to object-oriented models, now culminating in this agent-based approach [9]. Below we consider some of the fundamental properties of the agent approach, and distinguish classes of agent systems.

2.1 Agent Concepts in Different Disciplines

It has become acceptable to use the term "agent" in a dizzying array of approaches and applications:

Robots: The pedigree of agent concepts in technical fields probably begins with robots as the paradigmatic examples of autonomous agents. Increasingly, behavior-based robots are seen as even better examples, since their behaviors are emergent, and thus more autonomous, rather than pre-programmed [Rocha, pp. 2-3].

Information Systems: The use of agent concepts specifically arose in information systems, where agents are commonly thought of as independent actors usually acting as stand-ins for users in various negotiation tasks. Common are "helpers" agents. In help systems in single applications, agents are sometimes given anthropomorphic properties, as with "avatars" (Microsoft's ill-fated "Bob" is an example). In networked environments, independent agents are deployed to gather information for users autonomously. Recently economic negotiations are being used as agent-based applications [11, 26].

Software Engineering: Agent conepts have also quickly grown to be important in standard software engineering. Indeed, objects (as used in object-oriented engineering approaches) have many "agent-like" capabilities, especially encapsulation. It has become common to think of agents as "super-objects", combining encapsulation with autonomous process control and indpendent threading. Such ideas are the natural continuation of old-fashioned daemons and even DOS TSRs, which have agent-like capacities in terms of their asynchronous availability.

Artificial Life: The ALife research community also commonly describe their applications as "agent-based". Many ALife applications can best be described as "collective automata", where a relatively large collection of relatively simple state-determined systems are connected according to various complex temporal or topological schemes in order to demonstrate "emergent behavior".

We note that most of these models are effectively implementations of distributed dynamical systems with certain network topologies. An example is the wok of Kautz, Selman, and Milewski [23], who use a relatively simple Markov process distributed over a random graph to represent the referral pattern of expert knowledge over a network. While they use both agent and AI terminology, there is little in the work which requires or even uses these agent ideas explicitly.

Other ALife approaches embed their agents in virtual environments where the dynamical properties are coupled to interactions with environments and other agents. One example here is the work of Ackely and Littman [1], discussed in [Rocha pp. 11-15]. Another is that of Pepper and Smuts [28], who demonstrate the development of altruistic behavior, where it has not been found before, in virtue of coupling agent models in interaction with virtual environments.

Artificial Intelligence: In AI, complex actors with a great deal of on-board computational intelligence and planning ability are commonly and increasingly described as "agents". Of course the AI approach is radically different from the ALife schools. Recently Sloman and Logan [33] have provided a good discussion of the agent concept as used in AI. While they recognize a diversity of approaches, they emphasize the role of planning dedicated to solving specific tasks, and this is the orientation which dominates this agent approach.

An example is the work of Stilman [34], who deploys agent models of aircraft systems in a virtual reality of a small grid. It should be noted, however, that Stilman's agents actually have no autonomy, as their behavior in the grid is governed strictly by their deterministic (and off-board) planning systems.

Another example is that of Delgado *et al.* [8], who use a highly specific and complex fuzzy rule based learning and planning architecture among a small collection of agents for the purposes of approximating an analytic function.

Decision Theory: Finally, political scientists have also jumped on the agent bandwagon. In particular, in theories of how a group of individuals come to a collective choice, each individual in the group is represented as an agent. These researchers are mostly concerned specifically with the decision making capabilities of agents, either collectively or in groups, rather than other aspects of agents in general.

Examples include the work of Richards [31] [Rocha pp. 18-19], Wolpert [36], and their colleagues. We can see these approaches as the ultimate departure of the agent concept from its roots in robotics, in that pure decision making is considered divorced from any interaction with environments, either real or simulated.

2.2 Properties of Agent Systems

We are interested here in abstracting away from these disparate senses of agency and large collection of specific examples of agent systems, in order to discern first their common properties, and then their defining properties. First, we can see agent concepts clustering around a relatively small set of application types:

1. In information systems, to help with their simulation and engineering, as well as their user interfaces (for example the helper bots).
2. For the simulation of complex dynamical systems, as with most of the ALife applications.
3. In the simulation of natural systems such as organisms, humans, ecologies, economies, and societies.

Our interests extend to a mixture of 2 and 3, attempting to cast human organizations as collections of agents, but with aspects of dynamical systems more typical of collective automata.

Then, we find that agents have most or all of these properties:

Asynchronous: Agents act independently in time, commonly implementing some mecahnism for concurrency, parallelism, or independent control threads.

Interactive: Agents communicate and interact in a somewhat "social" manner, forming a collective entity through their interaction.

Mobile: Agents have some form of capacity for "movement", although this can have many different aspects: movement in a real, virtual, or simulated space, or movement of code or data among agents or between agents and their environments [24].

Distributed: Again, individual agents which make up agent systems are distributed in a real, virtual, or simulated space.

Random: Finally, some aspect of agent systems is non-deterministic. Either individual agents have non-deterministic behavior, or a large space of parameters of deterministic agents is searched through statistical trials.

2.3 Semiotic Agents: Autonomy and Freedom of Decision-Making

While the list of characteristics above is common to most agent models, they still do not capture the essential qualities which most people bring to the concept of "agency". These qualities are a kind of *independence*, the fact that the agent is doing something of and by itself. This refers to a kind of *self-control*, or, in a word, *autonomy*. From this property alone, all of the above follow.

Autonomy is an old concept, literally meaning "self-governing", from the Greek *auto* (self) and *nomos* (law).[1] Autonomy has connotations of both independence and separateness, and was originally used to refer to political organizations, as in national autonomy. Thus we recognize the following aspects of the concept of autonomy [20]:

Boundaries: Autonomy assumes that you can distinguish a particular domain over which the agent has control, and thus a boundary between what the agent controls (what it is autonomous over) and what it does not. But boundaries can exist with respect to many modalities, for example spatial, temporal, or functional boundaries.

Quantitative: Similarly, it is clear that autonomy admits to *degrees*, that something can be more or less autonomous, perhaps more in one of the modes mentioned above than in another.

Identity: The existence of a (perhaps partial) boundary implies some form of discreteness, and thus the ability to distinguish between that which is inside and outside the boundary. What is distinguished as inside thereby forms the identity of the agent.

Closure: Finally, everything we've said so far implies a form of (again, perhaps partial) closure, where those aspects of the world which are entrained within the boundary (identity) of the agent are then closed off from other interactions. As with boundaries, closure can take many forms, from physical (structural boundaries), causal (some form of encapsulation), or functional (closure of input/output mapping).

The list of properties above is actually quite familiar to us. First, these ideas are present either implicitly or explicitly at the foundations of systems theory [2, 4, 21]. Indeed, based on the above criteria there is very little to distinguish an "agent" from some general sense of "system".

Secondly, it parallels in some ways Holland's categorizations [13], as discussed in [Rocha pp. 3-4]. Boundaries and closures are all fundamental to the ability to aggregate (wether internally in terms of categorization or externally in terms of collective emergence); his tagging is essentialy our identity; etc. Rather, I would assert that Holland's list is actually rather close to a description of what would constitute a general systems approach to the delineation of various classes of meta-systems (multi-system systems). Indeed, it can also be claimed that other movements in computer science, for example the rise of object-oriented concepts, mirror the general systems approaches.

So what we need to do is approach a coherent sense of agent that will be distinguished not only from other software engineering senses (agents are not just subroutines or objects), but also from "objects" or "systems" in general (agents are not just systems). To do so, we focus on the concept of **autonomy with respect to action**. In other words, our concept of agent is a system (object)

[1] Note that this is not, as is commonly thought, etymologically related to self-reference, in that autonomy is not about self-naming, from the Greek *nomen* (name).

which has an inherent *freedom* to make *choices* or *decisions* over possible *actions*. This is what Rocha, from Aquinas, refers to as *election* [Rocha, p. 1].

We will call such agents *semiotic agents* to distinguish them from all others. We will discuss the use of the term "semiotic" in this context below. Therefore, following Rocha's discussion of dynamical coherence vs. incoherence [Rocha pp. 1-2], we can distinguish broadly semiotic from "dynamic" agents.

Dynamic:

- Possess functional or causal autonomy.
- Are dynamically coherent with their environments.
- Have input, output, and state.
- Allow dynamic self-organization (attractor behavior).
- Examples:
  - Physical systems following natural laws
  - Purely instinctual agents following natural propensities

Semiotic:

- Possess autonomy of action
- Dynamically incoherent with environment
- Also have memory.
- Examples:
  - Software agents of sufficient complexity
  - Organisms, people

Of course, this distinction is another expression of some common ideas in the literature. In particular, it is close to AI concepts of "reactive" vs. "deliberative" processes [33], and indeed, we would argue that all (deliberative) AI systems are semiotic in that sense. However, we are also motivated by the ALife and complex systems critique of AI, which allows for emergent phenomena through autonomy as opposed to external programming of elaborate internal models and planning mechanisms. Thus our goal is to construct semiotic agents which are sufficiently, but *minimally* sufficiently, complex to have autonomy of action.

3 The Semiotic Approach to Agent Modeling

3.1 Systems Foundations

Our perspective on semiotic agents is rooted in systems-theoretical foundations. In particular, we distinguish systems (and thereby agents) from their environments, and recognize the possibility of emergent behavior only from system- (agent-) environment interaction [32]. A single agent is thus considered related to a (perhaps virtual) physical environment (left side Fig. 1). We will call this its "absolute environment".

Of course we said agents come in collections, or "societies". In this case, the system-environment distinction can become complex. In particular, as shown on the right of Fig. 1, agents in multi-agent

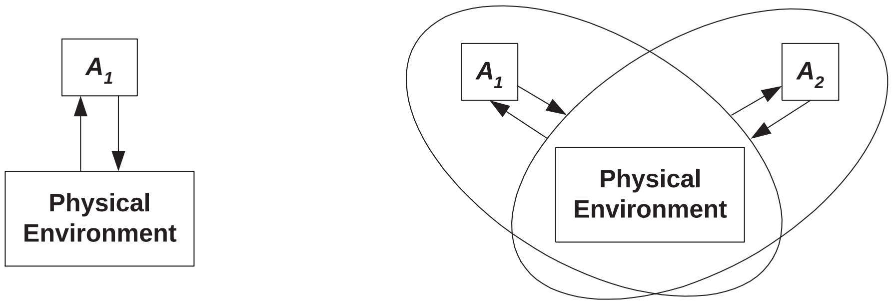

Figure 1: (Left) An agent interacting with its absolute environment. (Right) Agents interacting with their relative environments, also consisting of other agents.

systems interact with "relative environments" consisting of the "physical" environment *as well as all the other agents.*

In particular, if agents *do* have internal models of their environments, then these models must represent not only the absolute ("physical") environment, but also the other agents, and in particular the other agent's models. In this way, we arrive at various "reflexive" modeling strategies, for example as developed by von Foerster [35] or Lefebvre [25].

Considering collections of agents (agent systems) further, we can consider a variety of potential types:

Small collections of simple agents: In the limitting case, a few simple state-determined automata can be composed into higher-level automata. Similarly, simple robot interactions can frequently be modeled completely, although at a higher hierarchical level. In general, such systems have the possibility of analytical global descriptions.

Large collections of simple agents: This is the traditional ALife approach, where a large collection of simple state-determined automata are combined according to particular network or topological relations. The large collections allow for emergent phenomena at the level of dynamical attractors (a now-classical example is the work of Kauffman [22]). In the limitting case, this approaches statistical physics, and the approaches taken are similar, relying on experimentation and statistical descriptions.

Small collections of complex agents: This is the traditional AI approach, where a relatively small collection of highly intelligent agents interact cooperatively or competitively. Usually, both the environments, tasks, and planning strategies used are highly tuned to the application desired, and there is little empirical exploration of a variety of possible states, parameters, or initial conditions, let alone architectures.

Our goal is to aim solidly between the ALife and AI approaches, implementing agents which are relatively simple, and thus whose collections can have emergent properties, but with sufficient memory bases and uncertainty structures to allow for deliberative capabilities.

3.2 The Architecture of Semiotic Agents

The fundamental architecture we are proposing for semiotic agents is shown in Fig. 2. The system takes measurements from its environment, and constructs generalized “beliefs”: stored representations of the current state and memories of past states. There is also an internal representation of “desires”, namely potential goals states. A decision node decides among potential actions, which are then taken back into the environment. Finally, those actions interact with the dynamical processes in the environment, which then feed back to the agent in the form of future perceptions. In this way, the consequences of the agent decisions have an explicit impact on its future development.

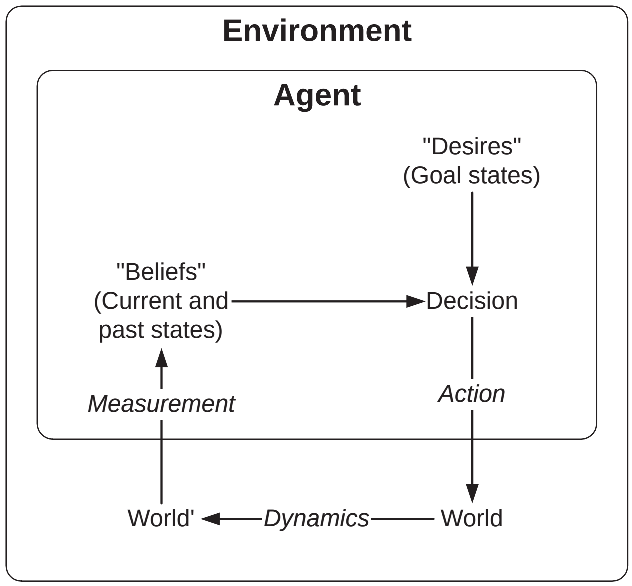

Figure 2: Architecture of a semiotic agent.

This architecture is based on the principles espoused above, in particular:

Generalized Control Architecture: The autonomy of the system is allowed by its manifestation of a closed causal relation with its environment, in particular a generalized control relation [29, 30]. Through this relation, the agent takes decisions to make its beliefs as “close” as possible to its desired state in order to reduce a generalized “error function”. Thus semiotic agents manifest a generalized negative feedback control relation.

System-Enviornment Distinction: The agent’s tight coupling with its environment across the system-environment boundary is absolutely essential. The consequences of agent decisions are always reflected, through the environmental dynamics, back to the agent in terms of more stored memory states.

Deliberation and Election: The autonomy of action necessary for semioic agents is allowed in virtue of the dynamical incoherence of the memory structure and the indepenent representation of the decision function.

Simplicity: Semiotic systems are not AI systems. Rather, beliefs and desires as relatively simple, non-propositional uncertainty structures.

3.3 Semiotics

So far we have not motivated the specific use of the term "semiotic" to describe our agents. What characterizes these systems is that they involve processes of perception, interpretation, decision, and action with their environments. The memory structures required by such dynamically incoherent systems fruther entails the presence of *representations* stored internally to the agent, in particular of measured states of affairs, goal states, and possible actions.

It is for this reason that we turn to semiotics, or the general science of signs and symbols. Originally a humanities sub-field of linguistics [6, 10], semiotics has come to become more prominent first in text and media analysis, and then in biology [7], computer engineering, and control engineering [27].

Semiotic processes involve the reference and interpretation of sign tokens maintained in coding relations with their interpretants. Thus semiotics in general is concerned with issues of sign typologies, digital/analog and symbolic/iconic representations, the "motivation" (intrinsic relations of sign to meaning) of signs, and mappings among representational systems.

Semiotics further decomposes semiotic relations along three axes:

Syntactic: Concerning relations among sign tokens, the production of new tokens, and the formal properties of symbol tokens as used in symbol systems.

Semantic: Concerning the interpretation of tokens by agents as standing for environmental observables.

Pragmatic: Concerning the repucusions of those interpretations for the agent in its environmental context, in other words, the purposes or goals of sign interpretation.

Semiotic relations are characterized by being **codings**, or in other words **contingent functional entailments**. In particular, they are entailments, meaning regularities of constraints in system relations; which are functional, meaning deterministic (equivalent to a mathematical function); and which are contingent, namely that other such functional entailments (coding relations) could have been possible [17]. This concept captures the arbitrary coding nature of symbol systems: the symbol and its referent share no properties in common except that the symbol refers to its referrent when interpreted by an agent acting within the constraints of the symbol system. These are contrasted with purely physical systems, which are characterized by necessary functional entailments. Note that this distinction roughly parallels that between dynamical coherence vs. incoherence.

A simplified example will serve to illustrate this point. Let O be a simple organism which lives near an oceanic thermocline with warm water above and cold water below. O acts as a semiotic control system in relation to the thermocline. Its perception is a single critical variable of temperature with states

$$X = \{+ = \text{too hot}, - = \text{too cold}, 0 = \text{just right}\},$$

and it has a single variable action with states

$$Y = \{u = \text{go up}, d = \text{go down}, n = \text{do nothing}\}.$$

There are $3^3 = 27$ possible functions $f: X \mapsto Y$, any of which the agent could invoke to make a decision to take a particular action. But only the three shown in Table 1 will result in stable negative feedback control. In all other cases, positive feedback, and not negative feedback, will

x	$f_1(x)$	$f_2(x)$	$f_3(x)$
$+$	d	d	d
$-$	u	u	u
0	n	d	u

Table 1: Functions sufficient for semiotic control.

result, with a corresponding runaway behavior: either the organisms will continue to ascend when warm, and thus boil; or descend when cold, and thus freeze. While any of the three will result in the organisms survival, f_1 is the best default selection, since it minimizes unnecessary action and results in smoother and faster control.

Note that there is no fundamental natural law of the universe which requires f to be selected according to the principles of negative feedback. Instead, this selection is *contingent on*, and *results from*, the process by which the system is *constructed.*

3.4 Consequences of the Semiotic Perspective

There are a number of important consequences which follow from the adoption of this semiotic approach to agent modeling.

3.4.1 Bounded Freedom on Decision Making

Perhaps the most fundamental is the recognition that semiotic agents operate in a context of *bounded freedom* on their decision-making capacities. We have emphasized that they have some freedom over decisions, otherwise they would not have autonomy of action. But on the other hand, they operate in contexts in which there are constraints from many sources. Recently, researchers have demonstrated that such constraints can be crucial in providing robustness and stability in multi-agent systems. These include:

Virtual Physics: Agents can be embedded in a virtual physical environment, whether simulating aspects of a real environment or a purely synthetic world. Decisions about actions are thereby constrained relative to the properties of these environments.

Gordon and Spears [12] have simulated distributed sensor grids exploiting an agent model interacting with an environment which manifests a certain limitted virtual physics. They have shown that they can achieve hexagonal or square grids based on the dynamics of the agent interactions with those "natural laws".

Similarly, Pepper and Smuts [28] have demonstrated the development of cooperative and altruistic behavior in simulated ecologies, but *only* when the environment had a rich enough "texture" of simulated vegetative diversity.

Communication: Agents can coordinate actions and learn about the physical environment through communication. Decisions about actions are thereby constrainted by the semiotic systems used to record, transmit, and interpret information.

Perhaps the best example here is the long-standing semiotic work of Edwin Hutchins [14, 15] (see also [Rocha p. 17]). In both real naval STOs and agent simulations of communication processes, Hutchins and his colleagues have demonstrated that the ability to create, manipulate, share, and interpret shared symbol tokens within sign systems is a necessary capability to develop robust organizational dynamics and collective cognition.

Shared Knowledge: Finally, decisions of agents may be constrained by a shared set of knowledge or beliefs, for example through a common biological evolution or cultural transmission (training or education). A cogent example here is the work of Diana Richards and her colleagues [31], discussed fully in [Rocha pp. 18-19].

3.4.2 Agent Dependence

Knowledge in these systems is necessary agent-relative, implying a kind of subjectivity, relativism, or constructivism. This follows from the fact that semiotics emphasizes the necessity for signs to be *interpreted* by agents in order to be meaningful: signs (symbols) never have meaning in and of themselves, but only as interpreted by an agent. Thus knowledge is local, and agents only have access to the world-as-perceived. There is then a dependence on measurable quantities, many of which are given from the construction of the agent.

3.4.3 Internal Models

So far, we have introduced semiotic agents as deliberative control systems involving internal representations of their environment, but not yet explicitly involving internal models. We have argued [19, 20, 21] that models and control are distinct, but canonical, examples of semiotic systems; and further that model-based or anticipatory control, where explicit predictions of the consequences of future actions are used to make decisions, are a necessarily more complex form of control than might be required. This is a difficult and deep point in the history of systems theory [5], and we will consider it further elsewhere.

3.4.4 Dynamical Opacity

We have argued that the freedom of choice which semiotic agents necessarily have is related to a form of nondeterminism. Indeed, it can be argued that a simple stochastic automata has many of the characteristics of a semiotic agent, and conversely a complex semiotic agent might appear to act as a simple stochastic process from an external perspective.

In other words, our common sense of election implies an entity making an intentional, deliberative choice among a set of possible actions. However, from an external perspective, it might not be possible to determine whether the system is acting with this sense of freedom or is simply a non-deterministic system. Moreover, a system may in fact be deterministic, but of such complexity that we can simply not identify its transfer function.

To a certain extent we are comfortable equivocating between an "ontological" perspective based on how the system "really" is, and an "epistemic" perspective which is only concerned with possible external models which can be constructed so as to describe systems. In this sense, we claim that semiotic agents have a form of **dynamical opacity**, in that they cannot be modeled as dynamical systems or simple state determined systems, even if they are in fact implementing some form of formalism. Of course, any deterministic system of sufficient complexity (and any chaotic system) can fall prey to this condition.

So our position is not that semiotic systems necessarily *are* deliberative systems, but rather instead that it is required that they be *modeled as* decision-making systems. Thus while we are aiming to construct agents with deliberative capacities, we are prepared to admit systems with simpler arhictectures as semiotic agents.

3.5 Summary: The Semiotic Approach to Agent Modeling

Here we summarize some of the important conclusions to draw about agent modeling from a semiotic perspective, before turning our attention to STOs in particular.

Environment: Perhaps the most important consideration is that simulated agents operate within environments which *have their own rules*, or their own "virtual physics".

Action Capabilities: Agents have action capabilities which must be considered *relative to those environments.*

Decision Capabilities: The possible decisions that agents can make must be considered *relative to those possible actions.* Thus we assert that pure decision models such as [31, 36] cannot fully realize the full emergent capabilities of agent communities in complex environments.

While we should not focus on a decision capability to the exclusion of a broader simulation, nevertheless this is clearly the most *important* component, and of most interest to this research project. Rocha [Rocha] has provided a thorough survey of classes of decision-making capabilities in agents, from encounters to strategies to learning and evolution. We will just summarize these here, in increasing order of complexity, as:

- Deterministic input/output state systems.
- Mutable transfer functions in terms of evolutionary (external selection) or adaptive (internal selection) processes.
- Finally, the use of culture as shared knowledge among agents to aid in agents decision-making.

Data: Data is seen as *information transmission among agents.*

Knowledge: Knowledge is seen as the *interpretation of data* by agents

Internal Structures: Can be characterized in terms of *state, memory, and decision* functions.

Communication: Among agents must be seen as *relative to the knowledge and internals* of the agent.

Control: Is seen as a form of *decentralized constraint over decision-making* in agents, potentially from many sources, including everything above.

4 Socio-Technical Organizations

We now move our considerations closer to the application area. There is currently a great need to bring computational, simulation, and information scientific tools to bear on the problem of representing and controlling complex systems which involve a great deal of interaction between human organizations and computer-based distributed information systems.

We call such systems **socio-technical organizations** (STOs), prime examples of which include:

- Hierarchical command organizations such as 911/Emergency Response Systems (911/ERS), search and resuce operations, and military organizations.

- Utility infrastructures such as power grids, traffic and transportation systems, gas pipelines, telecommunications systems, electronic markets, and the Internet.

The pressing needs are to assess the stability and vulnerabilities of STOs, and to protect their robustness against disruption in the event of destablizing forces, such as inherent dynamical instability, structural modification, or information disruption or disinformation, perhaps through deliberate attack or sabatouge. Our contention is that semiotic agent modeling approaches will be useful for simulating systems of this kind.

4.1 The Architecture of Socio-Technical Organizations

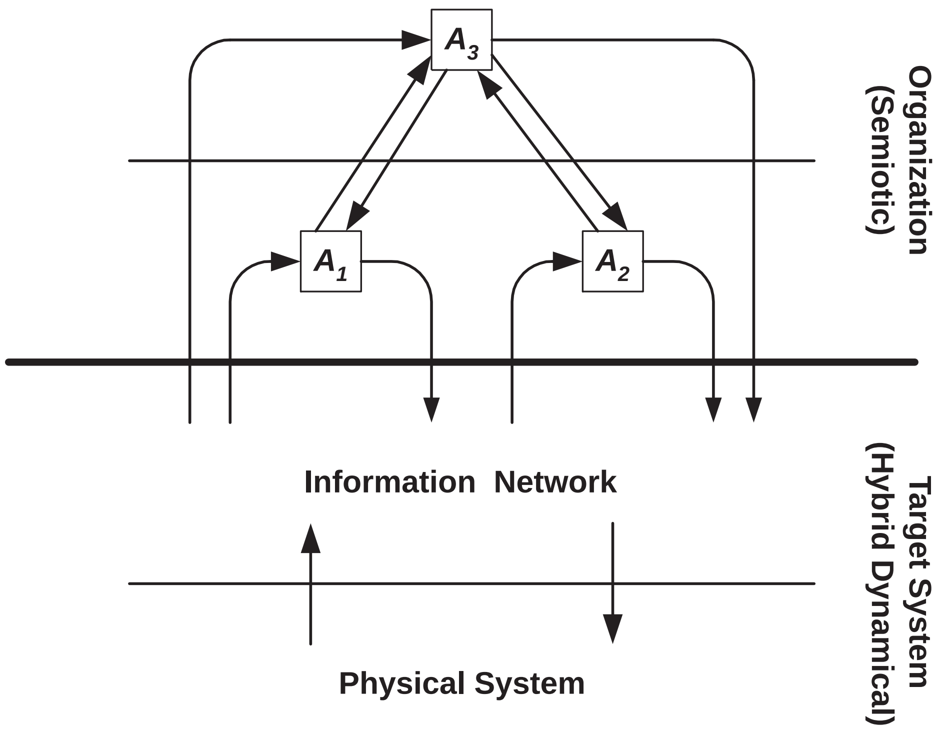

Figure 3: Socio-Technical Organizations.

STOs are characterized by a complex structure involving the hybrid interaction of physical systems with agent (human) organization. This involves, in order of increasing time scale, as shown in Fig. 3:

- At the lower level is a **target system**, which itself consists of two levels:
  - At the lowest level is a **physical system** which is deterministic (typically, and as we will assume here, a continuous dynamical system), involving the flow of physical objects or substances through a complex enviornment ("terrain").
  - Above that is an **information network** which is semi-automated, largely computer-based, and dependent on data acquisition, telemetry information, and control actions with the dynamical system.
- The target sytem acts as the environment to an **organization** of (human or computationl) agents or actors, which also has a complex structure:
  - At the lowest organizational level, **operators** are atomic units which interact in prescribed ways with the information network.

- At higher **supervisory** levels, supervisors can establish operational boundaries over lower or parallel systems, and alter system parameters.
- Ultimately, the highest organizationl levels involve the goals of the various corporate, military, and/or governmental organizations involved, including economic and political forces.

4.2 Distinctions Among Components

In any particular STO, the boundaries among these levels may be drawn very differently, or certain levels omitted. We distinguish the boundary between the target system and the organization by those components which *must* be considered as semiotic agents, and those which might not be. So based on our argument above concerning dynamical opacity, in general what distinguishes the target system is that it can be modeled as a deterministic, dynamical system, while the organization cannot. There are potentially a number of reasons for this, for example missing data about, or the computational complexity of, the organizational level.

In particular, a human, if sufficiently constrained by conditions in the environment or communication system, might be representable as a deterministic component of the target system; and conversely, a computer system of sufficient freedom and complexity might be considered part of the organization.

5 Agent Simulation of Structured Command Organizations

Our project in particular is intended to simulate the emergent decision structures in a 911/ERS system. We look at 911/ERS as an instance of a sub-class of STOs which we call **Structured Command Organizations** (SCOs). Here we describe SCOs and consider some special attributes of their simulation.

5.1 Structured Command Organizations

SCOs are characterized by a number of special properties:

- The organization contains a large number of units.
- The units are hierarchically organized, both for information flow upward and command flow downward.
- The lowest level units are individual humans, perhaps in vehicles.
- The organization must achieve a goal within a distinct time and within a physical environment.
- The environment may or may not contain other organizations with which the SCO interacts.

Examples of such systems include:

Disaster Response Systems

911/Emergency Response: When generally characterized, these encompass everything from routine police and fire calls, to disaster respone as mentioned above, to full-blown national emergencies requiring virtually military multi-agency response.

Search and Rescue Operations

Non-combatant Evacuation Operations (NEOs)

5.2 General Issues for SCO Simulations

Here we conider some general issues regarding the simulation not just of SCOs, but simulation in general.

5.2.1 Scenarios

It is presumed that any agent simulation developed for this project will have to have specified at least the following components:

Game Environment: The "game board" and its properties, the game pieces (semiotics agents), their capabilities, and their goal.

Agent Internals: Belief, desire, and decision structures consistent with the semiotic architecture.

Information Network: Communication channels, modalities, and capacities among agents.

Agent Organization: Organizational structures, largely initial.

These components will define the basic operating scenario for the simulation, and will presumably involve at least movement of a hierarchically structured collection of semiotic agents, if not other capabilities (e.g. carrying capacity, retrieval, or other specialized actions).

5.2.2 Fidelity vs. Abstraction

A central concern in any modeling effort is the amount of fidelity strived for, with a balance struck between fidelity and abstraction. In particular, high fidelity will require that we build in a lot of realistic details into the simulation, for example:

- Making the atomic units individual "nodes" rather than mid-level units or just arbitrary unit at some unspecified level.
- Faithfully representing the communication channels, modalities, capacities, etc. (e.g. voice and images), rather than a limited number of arbitrarily assigned low-cardinality channels capable of transmitting sentences in some limited formal language pertaining only to the game environment.
- Faithfully representing real unit capabilities for e.g. movement and carrying, rather than arbitrary or relative values like "fast" and "slow", or simple "point values" for carrying capacity.
- Using the real quantitative scaling relations among units (e.g. four precints in a police district) rather than arbitrary scalings (e.g. "n level 2s in a level 3").
- And most importantly, representing the actual command hierarchy structure as a tree rather than allowing any arbitrary structure like a general directed graph (see below).

In each of these cases, if we take the former course rather than the latter, not only do we give up formal simplicity and elegance in favor of fidelity, we also build in initial structure and correspondingly decrease any amount of emergent structure which might develop within the simulation. It is this emergent structure which we have been explicitly tasked to try to support and recognize.

Finally, the level of abstraction selected will also determine to a large extent the form of both initialization and validation of the model. In a highly faithful simulation, specific systems are instantiated to match the "real" system to a high degree, and similarly only a low degree of divergence between simulated and real data are tolerated as a result. On the other hand, the results achieved will not generalize to other systems at other scales.

5.3 Hierarchical Structures

Perhaps the most important single property which SCOs is their hierarchical organization. Considering an SCO as an ALife-type problem, the topology or network structure cannot be a simple flat "collective automata", but must reflect this "echelon" strcuture.

We consider an echelon structure as a generic hierarchy. A minimally simple example (three levels, two fanout at each node) is shown in Fig. 4 as a point of departure (see also [3] for a military example). In this structure, command flows down, while information flows up, down, and across to siblings.

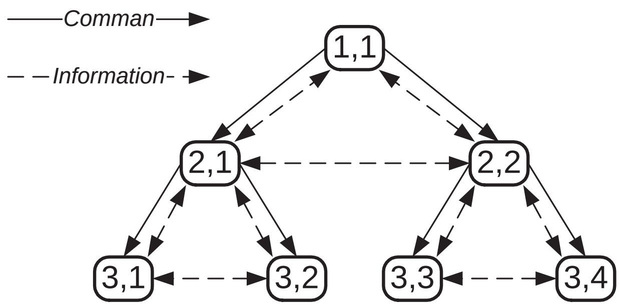

Figure 4: A generic three-level, two-fanout command hierarchy.

5.3.1 Alternate Hierarchical Structures

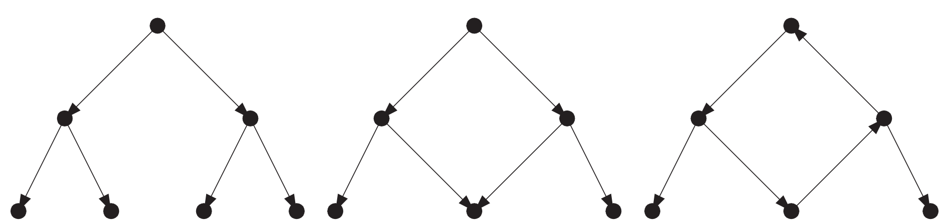

Figure 5: (Left) A strict hierarchy: tree. (Center) A loose hierarchy: DAG. (Right) A general graph: contains cycles.

One important consideration for us is to explore alternative structures. Examples of mathematically possible forms are shown in Fig. 5.

Trees: Classical command hierarchies are representable as mathematical trees, as shown in our "point of departure" example in Fig. 4. These are highly constrained structures, which we will call "strict hierarchies".

DAGs: Slightly less restrictive are general Directed Acyclic Graphs (DAGs). These are still hierarchical in the sense of having distinct levels, but these are "'loose hierarchies", in that a unit might have multiple parents (commanders) [16, 18].

General Graphs: Finally, even less restrictive are general directed graphs. The difference here is the possibility of cycles, wherein one can end up "commanding" oneself through a cyclic chain.

While SCOs do not intend to support either of the two looser structures, we understand that this might actually occur, at least for periods of time, for example near a boundary between two units. Loose hierarchical states may occur as transitions from one tree to another, or as stable states in and of themselves which require explicit representation, as shown in Fig. 6.

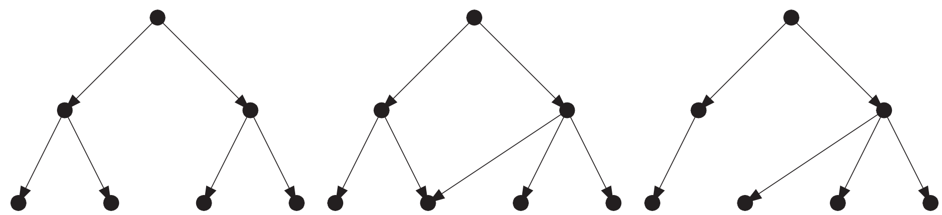

Figure 6: Loose hierarchical structures (DAGs) as transitions from one strict hierarchy to another.

Presumably an SCO would not tolerate such a situation, and strive to reconstruct itself into a new strict hierarchy if an unstable situation temporarily destroyed the structure. In this way what we're talking about are transitions from one hierarchy to another, rather than some kind of inherently non-hieararchical structure.

Finally, a simple theorem from discrete math says that you can "dangle" a tree from any node and produce another tree. Fig. 7 shows an example where the interior node b can be elevated to a "root" status. There has been speculation as to how it might be possible for such situations to arise under dynamic conditions, where due to local developments, a particular mid- or low-level unit encounters a crisis situation (for example, locating the rescue target) and then becomes the focus of resources. The issues here are rather subtle, involving the relation between command and control as a generalized form of constraint, which is considered a bit more below.

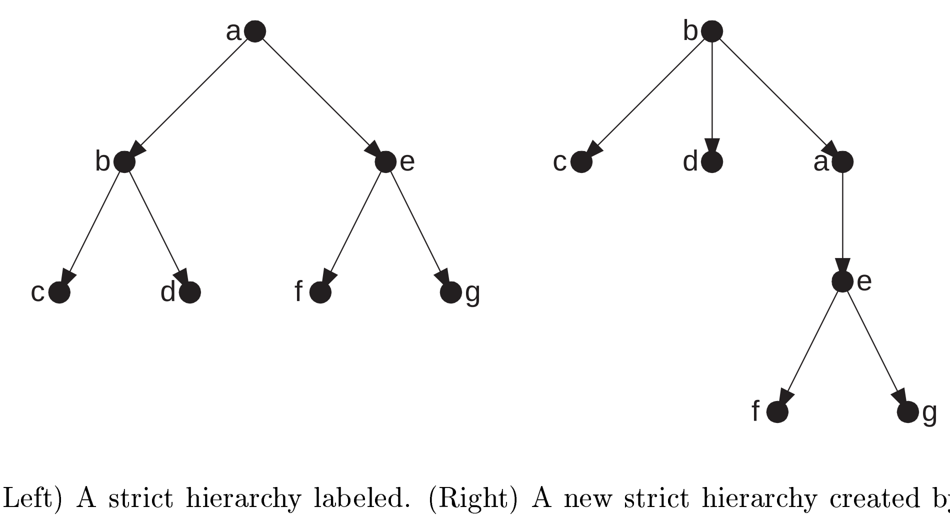

Figure 7: (Left) A strict hierarchy labeled. (Right) A new strict hierarchy created by elevating b to the root.

5.3.2 Structural vs. Functional Hierarchy

So far we've been discussing the echelon hierarchy. This is clearly a structural hierarchy of composition: a company consists of its platoons, and thus, for example, to move the company means to move its constituent platoons. There is a corresponding functional hierarchy embedded within the echelon structure: each unit at each level can move, gather information, take decisions, take actions, etc. Thus one can talk about e.g. the intelligence gathering hierarchy in parallel with the supply hierarchy.

Another consideration is that the structure shown in Fig. 4 is actually a two-way hierarchy, with command down and information up. We understand that units might actually participate in multiple independent hierarchies, reporting up to a different unit from which one takes command, as shown in Fig. 8.

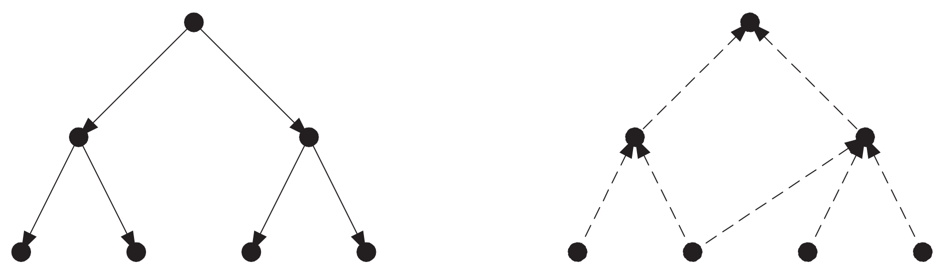

Figure 8: (Left) Command hierarchy. (Right) Reporting hierarchy for the same set of units.

5.4 Communication and Knowledge in the SCO Environment

Like command, information and knowledge is also distributed in a hierarchical manner in SCOs. And of course the agent organization is mediated by communication amongst its members.

We can identify "transmit" and "receive" as generic agent capabilities, and presumably the content could be represented as sentences in a simple formal language.

We can recognize multiple possible "channels" for the transmission of messages within the overall system:

Peer to Peer: "Horizontal" communication.

Superior to Subordinate: Transmission of new goals vertically down.

Subordinate to Superior: Transmission of new beliefs vertically up.

Environment to Agent: Measurement through the information network about anything in the virtual environment (here considered as the physical environment and the other SCOs within it) is probably best conceived of as communication from the environment to any agent.

Note that all of these except the last is anticipated by Fig. 4.

Also note that given a strict hierarchy echelon structure, any *emergent* structure will be very dependent on *horizontal* communication among peers within a unit (e.g. shift commanders within a police precint) rather then vertical communication among superiors and subordinates.

Possible message properties could include:

Cardinality: Continuous or discrete with various possible cardinalities.

Confidence: With a potentially hybrid uncertainty representation.

Meaning: Interpretation of a particular sentence by an agent.

Finally, we need to address the hierarchical nature of knowledge itself. In particular, along the echelon hierarchy a number of scalings interact:

Space: Higher level units operate over larger areas than lower.

Time: Higher level units operate over longer time frames than lower.

Scope: Higher level units operate with less detail than lower.

Presumably information is also structure according to these criteria, with high-level, abstract, large-area, and slow-to-update information at higher levels; and low-level, concrete, small-area, and fast-to-update information at lower levels.

Recall the emphasis that semiotics places on *data* vs. *knowledge*, in that knowledge is data *which is interpreted by an agent.* Thus even if there is a "broadcast" capability within an SCO, where there is some level of "universal data" available to commanders at all echelon levels, their *attention*, and thus the level of *interpretations* they make, and thus the *knowledge* they actually extract from the data stream, will be at radically different levels.

5.5 Command as Constraint

Finally, we have stated that in the semiotic approach we consider control as a distributed property entailed by various kinds of constraints from various sources. We must carefully consider how this idea interacts with a more traditional sense of "command".

Consider the following sources of constraint within SCOs:

Commander's Intent: Mission goals propogate down through the echelon hierarchy.

Operational Indoctrination: Generally there is a wide range of *shared knowledge* within an SCO, including training, maps, language, concepts, known responses and pre-made plans, etc.

Information Channels and Modalities: The fundamental bandwidths and modalities of the information network are a crucial form of constraint, and form the measurement input channels and action modalities for any STO.

Physical System: Finally, the "physical system" (as in the description of an STO from Sec. 4) places a number of constraints in the form of terrain, weather, equipment failures, target location, etc.

All of these factors end up *controlling* the actions of a particular unit at a particular echelon level, in the sense that they *constrain possible choices* and thus partially *limit the freedom* of the unit. Command itself is thereby just one more such constraint.

What about a situation where the non-command forms of constraint are actually more significant? Consider in particular the "dynamic situation" model introduced in Sec. 5.3.1 and illustrated in Fig. 7? If within a particular frame of time and space, a commander decides to let such a unit as b direct attention, command resources, and *effectively* constrain the freedom to decide of the other units around it, then this is, indeed, a form of control, whether or not we choose to recognize it as form of command.

References

[1] Ackley, DH and Littman, M: (1991) "Interaction Between Learning and Evolution", in: *Artificial Life II*, ed. C. Langton et al., pp. 487-509, Addison-Wesley, Reading MA

[2] Auger, Pierre: (1990) *Dynamics and Thermodynamics in Hierarchically Organized Systems*, Pergamon, Oxford,

[3] Baxter, Jeremy and Hepplewhite, Richard: (1999) "Agents in Tank Battle Simulations", *Communications of the ACM*, v. **42**:3, pp. 74-75

[4] Bunge, Mario: (1992) "System Boundary", *Int. J. of General Systems*, v. **20**, pp. 215-219

[5] Conant, Roger C and Ashby, Ross: (1970) "Every Good Regulator of a System Must Be a Model of that System", *Int. J. Systems Science*, v. **1**:2, pp. 89-97,

[6] Deely, John: (1990) *Basics of Semiotics*, Indiana UP, Bloomington IN

[7] Deely, John: (1992) "Semiotics and Biosemiotics: Are Sign-Science and Life-Science Coextensive?", in: *Biosemiotics: The Semiotic Web 1991*, ed. TA Sebeok, J Umiker-Sebeok, pp. 46-75, Mouton de Gruyter, Berlin/NY

[8] Delgado, Miguel; Gomez-Skarme, AF; and Marin-Blazqu, JG et al.: (1999) "A Multiagent Architecture for Fuzzy Modeling", *Int. J. Intelligent Systems*, v. **14**, pp. 305-329

[9] Dyer, Douglas E: (1999) "Multiagent Systems and DARPA", *Communications of the ACM*, v. **42**:3, pp. 53

[10] Eco, Umberto: (1986) *Semiotics and the Philosophy of Language*, Indiana U Press, Bloomfield

[11] Glushko, Robert; Tenenbaum, Jay M; and Meltzer, Bart: (1999) "An XML Framework for Agent-Based E-commerce", *Communications of the ACM*, v. **42**:3, pp. 106-114

[12] Gordon, Diana; Spears, William; and Sokolsky, O et al.: (1999) "Distributed Spatial Control, Global Monitoring and Steering of Mobile Physical Agents", in: *Proc. 1999 IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA 99)*

[13] Holland, John: (1995) *Hidden Order: How Adaptation Builds Complexity*, Addison-Wesley, Reading MA

[14] Hutchins, Edwin: (1996) *Cognition in the Wild*, MIT Press

[15] Hutchins, Edwin and Hazlehurst, B: (1991) "Learning in the Cultural Process", in: *Artificial Life II*, ed. C. Langton et al., MIT Press

[16] Joslyn, Cliff: (1991) "Hierarchy, Strict Hierarchy, and Generalized Information Theory", in: *Proc. ISSS 1991*, v. **1**, pp. 123-132

[17] Joslyn, Cliff: (1995) "Semantic Control Systems", *World Futures*, v. **45**:1-4, pp. 87-123

[18] Joslyn, Cliff: (1996) "Semantic Webs: A Cyberspatial Representational Form for Cybernetics", in: *Proc. 1996 European Conf. on Cybernetics and Systems Research*, v. **2**, pp. 905-910

[19] Joslyn, Cliff: (1997) "Semiotic Aspects of Control and Modeling Relations in Complex Systems", in: *Proc. Workshop on Control Mechanisms for Complex Systems*, ed. Michael Coombs, NMSU Press, Las Cruces

[20] Joslyn, Cliff: (1998) "Models, Controls, and Levels of Semiotic Autonomy", in: *Proc. 1998 Conference on Intelligent Systems*, ed. J. Albus and A. Meystel, pp. 747-752, IEEE, Gaithersburg MD

[21] Joslyn, Cliff: (1999) "Levels of Control and Closure in Complex Semiotic Systems", in: *7th Annual Washington Evolutionary Systems Conference*, ed. G. van de Vijver et al.

[22] Kauffman, Stuart A: (1990) "Requirement for Evolvability in Complex Systems: Orderly Dynamics and Frozen Components", in: *Complexity, Entropy and the Physics of Information*, ed. WH Zurek, pp. 151-192, Addison-Wesley, Redwood City

[23] Kautz, Henry; Selman, Bart; and Milewski, Al: (1996) "Agent Amplified Communication", in: *Proc. 13th National Conf. on Artificial Intelligence*, www.research.att.com/ kautz

[24] Lange, Danny B. and Oshima, Mitsuru: (1999) "Seven Good Reasons for Mobile Agents", *Communications of the ACM*, v. **42**:3, pp. 88-89

[25] Lefebvre, VA: (1998) "Sketch of Reflexive Game Theory", in: *Proc. Workshop on Multi-Reflexive Models of Agent Behavior*, pp. 1-41, Los Alamos NM

[26] Maes, Pattie; Guttman, Robert H; and Moukas, Alexandrou: (1999) "Agents That Buy and Sell", *Communications of the ACM*, v. **42**:3, pp. 81-91

[27] Meystel, Alex: (1996) "Intelligent Systems: A Semiotic Perspective", *Int. J. Intelligent Control and Systems*, v. **1**, pp. 31-57

[28] Pepper, John W and Smuts, Barbara B: (1999) "The Evolution of Cooperation in an Ecological Context: An Agent-Based Approach", in: *Dynamics of Human and Primate Societies: Agent-Based Modeling of Social and Spatial Processes*, ed. TA Kohler and GJ Gumerman, Oxford UP, New York

[29] Powers, WT: (1973) *Behavior, the Control of Perception*, Aldine, Chicago

[30] Powers, WT, ed.: (1989) *Living Control Systems*, CSG Press

[31] Richards, Diana; McKay, Brendan D; and Richards, Whitman A: (1998) "Collective Choice and Mutual Knowledge Structures", *Advances in Complex Systems*, v. **1**:2-3, pp. 221-236

[32] Rocha, Luis M and Joslyn, Cliff: (1998) "Simulations of Evolving Embodied Semiosis: Emergent Semantics in Artificial Environments", in: *Proc. 1998 Conf. on Virtual Worlds in Simulation*, pp. 233-238, Society Comp. Sim., San Diego

[33] Sloman, Aaron and Logan, Brian: (1999) "Building Cognitively Rich Agents Using the Sim_Agent Toolkit", *Communications of the ACM*, v. **42**:3, pp. 71-77

[34] Stilman, Boris: (1997) "Network Languages for Concurrent Multiagent Systems", *Computer Math. Applic.*, v. **34**:1, pp. 103-136

[35] von Förster, Heinz, ed.: (1981) *Observing Systems*, Intersystems, Seaside CA

[36] Wolpert, David; Wheeler, Kevin R; and Tumer, Kagan: (1999) "General Principles of Learning-Based Multi-Agent Systems", in: *Proc. 3rd Int. Conf. on Autonomous Agents*, http://ic.arc.nasa.gov/ic/people/kagan/pubs/agents99_bar.ps.gz