Smart Agents and Organizations of the Future

Kathleen M. Carley
Carnegie Mellon University

Direct all correspondence to:
Prof. Kathleen M. Carley
Dept. of Social and Decision Sciences
Carnegie Mellon University
Pittsburgh, PA 15143
email: carley+@centro.soar.cs.cmu.edu
fax: 1-412-268-6938
tel: 1-412-268-3225
url: http://sds.hss.cmu.edu/faculty/carley/carley.htm

This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-97-1-0037, NSF IRI9633 662, NSF KDI IIS-9980109, and the Pennsylvania Infrastructure Technology Alliance, a partnership of Carnegie Mellon, Lehigh University, and the Commonwealth of Pennsylvania's Department of Economic and Community Development. Additional support was provided by ICES (the Institute for Complex Engineered Systems) and CASOS – the center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University (http://www.ices.cmu.edu/casos). The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research, the National Science Foundation or the U.S. government.

Smart Agents and Organizations of the Future

As we move to the 21st century technologist point to the rapid changes in social and organizational activity that are expected to result from advances in computational technology. There can be little doubt that technology is altering organizations. Artificial agents such as WebBot, robots, and electronic shoppers are joining humans and organizations in the ranks of the smart agents that "work" in and among organizations. Computers are coming to control, or are involved in the operation of, everything from the office and home environment to routine purchases to strategic organizational decisions. As computers become embedded in every device, from pens to microwaves to walls, the spaces around us become intelligent (Nixon, Lacey and Dobson, 1999; Thomas and Gellersen, 2000). Intelligent spaces are characterized by the potential for ubiquitous access to and provision of information among potentially unbounded networks of agents (Kurzweil, 1988). Yet, we have little understanding of how to coordinate organizations in which humans and artificial agents work side-by-side, let alone how they work in these intelligent spaces.

The industrial revolution enabled organizations to increase in size, number of divisions (Etzioni, 1964; Fligstein, 1985), level of bureaucracy (Weber, 1947), and level of hierarchy (Blau and Scott, 1962). Information processing became key. Increasingly communication became organized so that orders and performance reports flowed down and information, decisions, and exceptions flowed up (March and Simon, 1958). Individual opportunity became based on networks of connections among jobs rather than patronage or nepotism (White, 1970; Yamaagaata, Yeh, Stewman and Dodge, 1997).

New technologies, both at the manufacturing and at the communication level, enabled certain organizational designs and affected what was adopted (Beniger, 1987; Aldrich & Mueller, 1982).

Today, information processing, communication, and knowledge management became key. Changes in computational power, telecommunications, and information processing are affecting when, where and how work is done (DiMartino and Wirth, 1990; Sproull and Kiesler, 1991). Further changes in agriculture, manufacturing, transportation and technology are leading to the emergence of an increasingly mobile population and knowledge intensive organizations. New organizational designs are emerging such as network organizations (Nohira and Eccles, 1992; Miles and Snow, 1995) and virtual organizations (Lipnack and Stamps, 1997). In these new organizations, even though information processing is key (Tushman and Nadler, 1978), communication is not constrained to be vertical (Contractor and Eisenberg, 1990). Organizational design becomes a strategic exercise in establishing and managing these relations (Burton and Obel, 1998) Rather, the network of connections within and among organizations act to constrain and enable the flow of goods, services, agents and information.

Advances in engineering and computer science suggest further changes will be forthcoming in organizations as the population of smart agents in organizations expands and the space becomes intelligent (Carley, 1999a). This chapter explores the potential effect of such changes on organizations. We begin by exploring the nature of smart agents and organizations as computational systems. The argument is set forward that the space which organizations occupy will become intelligent and individual's infospheres will expand. Within this space, search is likely to become the dominant task. Within this environment, the ecology of networks will constrain and enable all behavior. The ultimate question that social and organizational theorists will need to address is what happens to these networks, to organizational performance, and to organizational design when artificial smart agents begin to populate these networks?

At least four paradigms in organization science speak to the potential impact of smart agents on new organizational form – structuralism, contingency theory, information processing theory, and social networks. Work on organizational design suggests that different architectures influence performance and there is no one right organizational design for all tasks (Mintzberg, 1983, Burton and Obel, 1984). However, independent of the task, the organizational form can be characterized in terms of networks (Nohra and Eccles, 1992) such that the linkages among agents influence both agent behavior (Krackhardt and Kilduff,1994) and organizational performance (Baum and Oliver, 1991). Further, changing technology results in alterations of traditional structures by altering the networks to produce new organizational forms (Powell, 1990). Work on information processing (Cyert and March, 1963) demonstrates that it is the limits to agents information processing capabilities that affect organizational outcomes and that taking such limitations into account leads to more accurate prediction of organizational performance(March and Simon, 1958). This work also demonstrates that there is an interaction between knowledge (e.g., training, what agents know, and their information processing capabilities) and structure in effecting organizational performance (Masuch, M. and P. LaPotin, 1989; Carley, Prietula & Lin, 1998). However, little of this work speaks directly to the role of artificial smart agents. An exception here is some of the work on transactive memory which suggests that storing individual knowledge about who knows what in databases may have the same performance enhancing effects as when known directly by humans (Moreland, Argote and Krishnan, 1996). Another exception is the work on information flow which suggests that artificial smart agents change both the topology of the underlying networks, speeds the diffusion of information, and yet may maintain or exacerbate information inequalities (Kaufer and Carley, 1993; Carley, 1996). Collectively, this work leads to the conclusion that networks, cognition and the interaction among the two effect organizational performance. While this suggests that smart agents will also effect performance, it provides little guidance for the nature of those effects.

For example, key work on organizational learning has looked at the issue of search but ignored the fact that the underlying network within and among organizations constrains that search. Organizational researchers have long recognized the importance of search as a strategic tool for organizational adaptation (Levinthal and March, 1981; March, 1996). Search is typically characterized as unconstrained and as taking resources from knowledge utilization in that agents are cognitively limited and so can do only one of the two information processing actions at a time. Hence strong organizational performance is seen to require both such search-based exploration and the utilization of known information (exploitation) (March, 1996). In contrast, the social network tradition argues that networks in which agents are embedded constrains and enables their search (Contractor & Eisenberg, 1990; Carley, 1999b). The examination of smart agents requires that the information processing and the network views be melded into a view of agents as not only boundedly rational (cognitively and structurally) but also in which their ability to search is function of their position in both the social network, the knowledge network, and their information processing capabilities. As such search is not traded for knowledge utilization. Further, in intelligent spaces search is being conducted by smart agents, some of whom are artificial, who themselves are able to learn and where the direction of that search is enabled and constrained by the underlying networks. A consequence is that as the density of these networks change and as different types of agents populate these networks the time scarcity and competition among ideas begin to determine organizational outcomes. As the networks expand in both agents and ideas search effectively slows and organizational outcomes become a function more of order of learning.

Organizational design is a complex system in which a large number of factors interact in non-linear ways to effect performance. Moreover, it is a dynamic system changing humans learn, as goals change, and so on. The presence of artificial smart agents in organizations adds further complexity by enabling greater quantities of information to be stored, meta-knowledge to be created, artificial agents to act on behalf of humans, more knowledge to be created, and so on. The effects of these changes are again non-linear.

A valuable approach for studying such complex non-linear dynamic systems is computational modeling. Within organization theory, a new perspective has emerged which blends the information processing tradition, the social network tradition with a more veridical approach to cognition than bounded rationality. This perspective is known as computational organization theory (Samuelson, 2000).

Computational organization science is a new perspective on organizations and groups that has emerged in the past decade in response to the need to understand, predict, and manage organizational change include change that is motivated by changing technology

(Carley & Gasser, 1999). In this chapter, a computational organization theory perspective has been taken to explore the impact of artificial agents and intelligent spaces on organizational change. Agent based modelsare used to enable theory building (Epstein and Axtell, 1997). In an agent-based model, each actor is modeled as an independent information-processing agent with a set of knowledge and potential actions. In this paper the agents have the ability to learn, to interact, and so to realize a dynamic social network. The dynamic worlds that can be explored using this approach enable the researcher to address fundamental social science and organizational questions. As experimental testbeds these models provide an environment in which researchers can explore and learn the effects of complex relations (Lant, 1994) and generate hypotheses that can be tested in other settings (Carley, 1999c).

Thus, in this chapter, initial insights into the potential impact of smart agents on organizations will be provided by doing a computational analysis. Using a computational model inspired by a network based approach to understanding organizations, a series of illustrative virtual experiments are run. These virtual experiments are directed at exploring the impact of moving into an intelligent space on performance. The results provide us with better insight into the in which organizations might behave in intelligent spaces.

Smart Agents

Agents are intelligent if, in order to respond to a stimulus, they must engage in cognitive activity acting upon a body of information. One characteristics of cognitive activity is that it takes longer than programmed reflexes (Newell, 1990). Agents are adaptive if they change their behavior in response to changes in information. Agents are computational if they have the ability to do any of the following: acquire, process, store, interpret, or communicate information and the connections among pieces of information. Smart agents are agents that are intelligent, adaptive, and computational. Human beings are the canonical smart agents. However, many other smart agents exist that differ from humans in the degree and/or type of intelligence, adaptivity, and computation they exhibit (Carley and Kaufer, 1994; Carley, 1999a). Such smart agents are both real, such as dolphins, and artificial, such as electronic personal shoppers, automated email answering and sorting systems, WebBot, robots, and avatars are all examples of smart agents.

These artificial smart agents exist, at least in demo versions, today. Recent work in a number of areas, including that on robots (Thrun, 1996), avatars (Benford et al., 1997), intelligent agents (Weiss, 1999) and demonstrates the viability of artificial smart agents as an entire new class of organizational agents. Artificial smart agents are capable of working and communicating within and among organizations, on their own or with modest human intervention. As these agents take their place in organizations, new forms of coordination, new organizational designs are likely to arise.

Simulation based decision aids such as the Virtual Design Team (Jin and Levitt, 1996; Levitt, et al., 1994) and ORGAHEAD (Carley and Svoboda, 1996; Carley and Lee, 1998) employ smart agents to create more realistic environments for examining group and organizational behavior. Such multi-agent simulations enable the comparison of existing and new forms of organizing for collections of smart agents – both human and artificial.

The similarities between humans and artificial agents suggest that the difference among types of smart agents is often a matter of degree, both quantitative (the amount of knowledge of the environment and the number of capabilities) and qualitative (the aspects of the environment that the agent attends to and the types of capabilities). Carley and Newell (1994) define a knowledge/capability space for characterizing the features of agents. As agents increase in the amount and type of knowledge that they attend to, increasingly considering real time situations, multiple agents, multiple goals, and historical situations the variety and type of responses available to them widens. As agents move from cognitively completely capable, the omniscient agent, to increasingly constrained agents, i.e. from the rational actor, to the boundedly rational actor, to the cognitive actor to the emotional cognitive actor they increase in their need for diverse actions. The omniscient agent has no need to collect information as it knows everything. As we move in this classification scheme from the omniscient agent in an environment without space, time, social, historical or cultural constraints to the emotional cognitive agent in the everyday space we are accustomed to, the agents become increasingly human like, until the model social agent is reached. Using this scheme, various types of agents can be classified and compared in terms of their capabilities and every behavior related to the minimally capable agent needed to generate that action. Further, every position in this space can be operationalized as a computational model of an agent. An implication of this view is that computational models can be used to examine the relative impact of different types of agents (human and artificial) on organizational behavior. Another implication is that recognizably different social and organizational behaviors will emerge from the model as the agent characteristics change.

The Nature of Organizations

Organizations, like human agents or simple artificial agents such as avatars, are computational systems. As noted by Carley (1999b) "[a]ny entity composed of intelligent, adaptive, and computational agents is also an intelligent, adaptive, and computational agent." Thus organizations are also smart agents; but, unlike the individual agents we have been discussing they are synthetic. A synthetic agent is an agent synthesized out of multiple sub agents connected by a plethora of networks. Organizations exist within, and are defined by, an ecology of networks. As we move into the future, the behavior of the organization as an entity as well as the behavior of the agents within it will be affected by the movement to more intelligent spaces. As will be seen, the effect of intelligent spaces will be to alter the size and complexity of the underlying networks — interaction, knowledge and information. This is true whether those networks are among agents at the organizational or individual level.

An ecology of networks

A variety of networks exist within and among organizations. The four key corporate entities — agents, knowledge, tasks and organizations — define a set of networks (see table 1). The focus here is on three of these — the interaction network, the knowledge network, and the information network.[i] Various aspects of organizations can be characterized in terms of these networks. For example, structure (such as the authority structure or the communication structure) is defined in terms of the interaction network, culture in terms of the knowledge network, and the potential data in terms of the information network. Properties of the organization can be measured in terms of any one of these networks or the collection of them.

Table 1. Networks of agents, knowledge, tasks and organizations

	Agents	Knowledge	Tasks	Organizations
Agents	**Interaction Network** *Who knows who* Structure	**Knowledge Network** *Who knows what* Culture	**Assignment Network** *Who is assigned to what* Jobs	**Employment Network** *Who works where* Demography
Knowledge		**Information Network** *What informs what* Data	**Requirements Network** *What is needed to do what* Needs	**Competency Network** *What knowledge is where* Culture
Tasks			**Precedence Network** *What needs to be done before what* Operations	**Industrial Network** *What tasks are done where* Niche
Organizations				**Inter-organizational Network** *Which organizations work with which* Alliances

For humans, as boundedly rational (Simon, 1955; 1956) or cognitive agents (Carley and Newell, 1994; Carley and Prietula, 1994) their decision making ability, actions, and performance hinges on their extant knowledge, social position, and procedures and abilities to manage and traverse these networks. Network management involves being able to search for relevant people and knowledge, dynamically generate and evaluate the value/capability of groups of people and/or knowledge that are networked together to achieve some goal, and asses the vulnerability of the system to various types of dysfunctionalities (such as loss of personnel or knowledge), and manage change in these networks. For humans, the networks that people operate on, and in, serve to constrain and enable further action and affect the efficiencies of such actions (Burt, 1992). Similarly, for artificial agents, being able to traverse the digitized version of these networks enables machine comprehension (Bookman, 1994). For example, WebBots that serve as personal shoppers are more intelligent if they are more able to navigate through the links between sites on the web.

A change in any one of these three networks can potentially result in a cascade of changes in the others. For example, when individuals learn something new (by interacting with someone in their interaction network) that evokes a change in the knowledge network which can result in a change in the interaction network (Carley, 1991). As another example, when new personnel are hired they may bring new knowledge with them. As current personnel leave, the available knowledge may be depleted.

Managing these changes is the key to knowledge management. Information technology has the potential to affect this meta-network in several ways. First, it can affect the number and types of nodes in these networks; i.e. with the advent of new technology comes new agents, new knowledge, and new connections among knowledge. Second, information technology has the potential to alter the way changes occur and their impact. For example, some suggest that holding data in databases, and knowledge systems like lotus notes provides organizations with the means to decouple personnel turnover and change in the knowledge network

What is an intelligent space?

Intelligent spaces are physical spaces where access to other agents (human and artificial) is ubiquitous, the scale in terms of number of agents and amount of information is large, cognition is distributed, and computers are often invisible. The physical world in which people work and go about their daily activities is becoming increasingly intelligent. An increasing number of the objects that surround us (such as microwaves, VCRs, computers, answering machines, personal digital assistants, cell phones and security systems) have some level of intelligence; i.e., these devices are able to communicate, access, store, provide and/or process information.

Ubiquitous access means that technology will exist to enable all agents to access or provide information wherever, whenever, and to whomever it is useful thus remotely enabling other agents to act. Whether agents can exercise this ability will depend on the norms, incentives, privacy regulations, and security measures adopted by the group, organization, or society. In terms of scale, huge quantities of information will be automatically collected and stored and processed by a potentially ever-increasing number of agents. Information, access to information, and information processing and communication capabilities (i.e., intelligence and cognition) will be distributed across agents, time, space, physical devices and communication media (Hutchins, 1991; 1995). As computers are miniaturized, made more reliable, and increased in power and storage capability, we can expect more devices to become intelligent. This increases the number of agents, but it also starts making computers invisible. A further aspect of invisibility is that the interface between the digital world and the analog world will become seamless. For example, speech recognition and synthesis software, automatic transcription software, face recognition software, all enable a more seamless interface between the digital and analog world.

As spaces become intelligent there will be unprecedented increases in the size and complexity of the interaction and knowledge networks in which people (and other agents) are embedded and the size and mobility of their infospheres. The term infosphere refers to the collection of remote instruments, appliances, computational resources (all of which may be artificial agents), as well as the agents (human and artificial) and information made accessible to a person by these systems from a person's working environment, such as the desk and office or the bridge of a ship. All agents have an infosphere; however, the size of that sphere may vary as he agents change physical location. The knowledge available in these infospheres includes what agents know, who they know, and what they know how to access. For humans, the size of their infosphere is largely determined by the type of immediately accessible technology. Thus, your infosphere generally becomes smaller as you move from your office, to your car, to the hallway, to a remote mountaintop.

As spaces become intelligent, we expect two things to happen. First, infospheres will become larger. Indeed, there may be an increase in the complexity in individual's infospheres and the associated interaction, knowledge, and information networks well beyond people's ability to manage and monitor this space. Second, infospheres will become mobile. Thus, as the agents move from office to mountain top infospheres will degrade by choice, rather than access to technology. Moreover, technological change may lead to non-linear rates of change in these networks. For example, when one-to-one communication exists, even if every agent learns something new each time, the maximum number of new links in the knowledge network each time is N – the number of agents. In contrast, technologies which enable simultaneous many-to-many communication makes it possible for the maximum number of new links in the knowledge network to grow by N*(N-1). Technological change also may lead to fundamentally different structures (Barley, 1990; Kaufer and Carley, 1994). For example, databases enable teams to reach consensus by interacting with the database and so sharing knowledge off-line rather than reaching a shared understanding through direct interaction. Most organizational theory does not consider the effect of infospheres on organizational performance. In contrast work on technology suggests that new technologies, by altering the infospheres will enable more outwork, more use of temporary employees, potentially better decisions, but mixed effects on culture (Harris, 1994; Kiesler, 1996; Worthington, 1997).

As intelligent spaces alter the infospheres, the networks in which people are embedded, and those to which they have access to, are likely to respond dynamically and become potentially unbounded. The theory of bounded rationality suggests that limitations on humans determine the level of performance the organization can achieve. As these boundaries are eliminated then performance should improve. It is important to recognize that technology does not eliminate boundaries but moves them; i.e., in intelligent spaces cognitive, social and institutional barriers will still exist. In 1956, Simon noted that (p. 130) the agent makes decisions using knowledge, which includes simplifications of reality that "may depend not only on the characteristics — sensory, neural, and other — of the organism, but equally upon the structure of the environment." Artificial smart agents are not going to obviate such limitations in the human, for all that they may provide external access to information and agents. Moreover, these artificial agents themselves will also be limited both “cognitively” in their information processing capabilities and “structurally” by their position in the social network. Further, it is reasonably safe to assume that coordination and communication will still center on knowing who knows who (the interaction network) and who knows what (the knowledge network). Agents will still act as gatekeepers. However, rather than affecting who can link to whom, they will limit who does link to whom by the way in which information is located and provided. As a trivial example, currently on the web, which site a user accesses is affected by which search engine is used and the way that engine prioritizes the located sites.

How might smart agents affect organizations?

Stories about the impact of technology on organizations abound. Artificial smart agents are expected to have a number of interesting characteristics. Three such characteristics that are important in the organizational setting are: boundary breaking, communication extension, and storage (Sproull and Kiesler, 1991). These three impacts are critical as they affect the fundamental information processing capabilities of the organization.

Artificial smart agents as boundary breakers

Within and among organizations boundaries exist. For many network organizations, the boundary between organizations, between what is “internal” and what is “external” has virtually disappeared (Sproull and Kiesler, 1991; Miles and Snow, 1995). Smart agents, will break still more boundaries. In particular, they will make permeable the boundaries surrounding people, tasks and resources. For example, avatars can act on behalf of a human to schedule appointments or answer routine questions. Thus, making the person effectively more available than otherwise to both organizational and non-organizational members.

Smart agents as communicators

Smart agents can communicate. The import of this, is that within a community of humans and artificial agents, these communications will affect what is learned. Thus the truth, accuracy, and frequency of the communications sent by artificial agents can alter organizational performance. For example, just having a database increases the effective number of agents by 1 and provides greater access to information, beyond that provided by just people. If each person in the organization has a personal avatar that increases the number of agents to twice the number of people.

Smart agents as storage

Smart agents store information. The import of this is that within a community of humans and artificial agents, the ability of artificial agents to store information can alter the likelihood that the stored knowledge is re-communicated. Such stored information can potentially alter the group knowledge or shared mental model. This in turn can alter organizational performance. For example, databases as repositories make information available whether or not the individual who provided the information in the first place remains with the company.

Coordination

How do we coordinate organizations when both human and artificial agents are present? How will the presence of these artificial agents alter the rate of organizational change? Organizational change occurs in a number of ways, ranging from internal changes to the culture to changes that affect overall performance in the market. Three key indicators of organizational change are information diffusion (Carley, 1991), consensus (Carley, 1985) and task accuracy (Carley and Svoboda, 1996). One of the most commonly attributed effects of smart agents is that they will enable information to diffuse faster. Additionally, since more information is expected to reach more people faster consensus is expected to occur faster. Finally, smart agents are expected to enable greater accuracy as they enable the analysis of more information.

We now examine whether smart agents can effect organizational change by altering the absolute and relative rates of information diffusion, consensus formation and task accuracy. Using CONSTRUCT-O (Carley, 1990; 1991, 1995, 1999b) a virtual experiment is run and the results evaluated to explore the impact of artificial agents and intelligent spaces on in effecting organizational change. Changes in the absolute and relative rates of information diffusion, consensus formation and task accuracy are examined.

Three aspects of intelligent spaces are explored – changes in boundary spanning, changes in communication, changes in storage. How might the movement to intelligent spaces impact boundary spanning, communication and storage? First, as noted, there will be effects due to scale. That is, there will be more agents and more information to which people will have access. To capture this effect a static equilibrium analysis is used, in which the effect of access to more agents is captured by examining worlds with an increasing number of agents. Similarly access to more information is captured by examining worlds varying in the overall amount of knowledge available.

Second, there will be new types of agents – such as databases and avatars. These agents will differ from humans by having different information processing capabilities. One of the key capabilities of databases is that they do not forget, and can store immense amount of information beyond the tenure of or life span of individuals. One of the key capabilities of avatars is that unlike their human counterparts, they are always available for interaction. In this analysis, we differentiate these agents in terms of the following information processing capabilities: initiation, sending, receiving (learning), availability (number they can receive from at once), and amount of initial knowledge. The differences in these agents are summarized in table 2. Since knowledge is modeled as a bit string the amount of knowledge is just the percentage of those bits initially known by the individuals. What knowledge the individuals know initially is randomly distributed over the bit string.

Availability is of course much more complicated than as characterized in this table. Humans, even without technology can interact one to many. Databases may be locked for data entry to one to one for concurrency control but can often be simultaneously searched by many. And so on. However, in this virtual experiment availability for databases and humans is limited to one-to-one. There are two reasons for this. First, this is the predominant mode of interaction for both types of agents. Second, allowing only the avatars to be one-to-many aids analysis so that the impact of that information-processing feature can be clearly distinguished.

Table 2. Characteristics of Agents					
Agent	Initiate	Send	Receive (Learn)	Availability	Amount of Knowledge
Humans	Yes	Yes	Yes	Only to one other at a time	50%
Databases	No	Yes	Yes	Only to one other at a time	5%
Avatars	Yes	Yes	No:	To many at a time	10%

There are, of course, other characteristics of the societies populated by these agents. In any society, at any given time, there is a set of available information; i.e., the union of the information known by all agents in the society. This information is available in the sense that, that information by virtue of being known by at least one member of the society is potentially available to all. All human agents know 50% of the available information chosen randomly from the set of available information. When a database is present, that database is initially set up containing little information (5% of the available information chosen randomly from the available information). Whereas, each avatar knows 10% of the information available in the society. For each avatar, its information is a subset of the information initially known by its human counterpart. Each avatar knows 10% of the overall knowledge but this knowledge is 1/5th of the information known by its human counterpart. Which of the information known by its human the avatar knows is chosen randomly. One of the most fundamental findings in sociology is the tendency of individuals to interact more with those to whom they are more similar, homophilly (McPherson and Smith-Lovin 1987; Carley 1991). Reasons for this are many including ease of communication, shared understandings, and comfort. Humans, and in this model all agents, are information seeking but given two possible interaction partners prefer the partner with whom they have more in common. A consequence is that all agents (humans and artificial) behave in a very boundedly rational fashion. In particular, the following human agent behaviors emerge:

1. Initially they act as though there is a small incentive to contribute to the database; i.e., they will contribute if they have no one else to interact with.
2. Experts, those with much information, will contribute to the database rather than repeating themselves to others.
3. As more individuals contribute to the database novel information, i.e., information that is not already in the database, others will follow.

One can envision a large number of other possible types of agents such as referential databases, books, and personal shoppers. The agents examined (humans, databases and avatars) were chosen because they have distinct information processing capabilities, corresponds to a type of agent currently existing, and are likely to continue to exist in the future. Thus, the selected agents are likely to play a role in the digital economy and in transforming current workspaces into intelligent spaces.

The simulation is run for a number of time periods until quiescence is reached; i.e., no new information is being communicated. A time period is a communication-learningrepositioning cycle in which all agents find a communication partner (that might be themselves), send and/or receive a bit of information, learn new information sent to them, and on the basis of their new total knowledge change their propensities for interacting with others.

Information diffusion is measured as the number of time periods until all human agents in the organization know a randomly selected piece of information. The higher this number, the longer it takes on average for information to diffuse to any particular individual. Consensus is measured as the number of time periods until the maximum number of people who are ever going to agree, first agree. It is possible that, given a task, not all people will agree. At some point, a maximum is reached in the number of people who agree. The time at which this maximum is reached is the time used for consensus. Two individuals are said to agree if they vote the same given the task. Since individuals learn their propensity to agree on the same tasks changes over time. Individuals can go in and out of consensus. Task accuracy is measured as the percentage of problems in a single time period that are solved correctly by the organization. An organization solves a problem correctly if the majority of people's vote is accurate. Sustained accuracy is measured as the number of time periods until the organization's accuracy has stabilized. At some point, a maximum is reached in accuracy. The time at which this maximum is first reached is the time used for sustained accuracy.

To see how smart agents might transform the workplace in terms of performance, information diffusion and consensus a virtual experiment was run in which the number of human agents, amount of knowable information, and the types of agents were varied. Variations examined are summarized in table 3. Three categories of each variable were utilized resulting in 27 virtual worlds, where a virtual world is a simulated society with a specific number of agents of each type with specific levels and distributions of knowledge. Each of these worlds is then simulated 10 times. Results reported are the ensemble average of these different runs. Note for clarity of results, worlds with all three types of agents (humans, databases and avatars) simultaneously present were not run.

As previously noted, knowledge is modeled as a bit string. Here worlds differ in the length of that bit string. Which of the bits the agents initially know is determinedly randomly. The percentage of the bits known is specified in Table 2. Humans and databases can learn; i.e., over time the number of buts that they know can increase. For humans and databases, when they interact with another human they receive a single bit of information from that source. If they do not already know this information they learn it. Humans can in a similar fashion learn from databases. In stark contrast, avatars are modeled as having a set of knowledge (a sample of what their associated human knows) and this knowledge does not change over time; thus, avatars cannot learn. Clearly this is a simplification of reality as the learning always occurs if it can and it is perfect. Future work should consider errors in this process.

The task being used is the binary classification task (Carley, 1992). Each time period the organization is faced with 25 tasks. Each task is of the form — decide if in this string there are more 1's or 0's. Each agent has access only to those bits of the task that correspond to information they know. Each agent decides that there are more 1's/0's if in the set of bits that the agent is looking at (for the information they know) there are more 1's/0's. This decision can be thought of as the agent's vote.

Table 3. Synopsis of Virtual Experiment

Variable	Description	Values
Population [†]	number of human agents	10, 20, 50
Knowledge	number of pieces of available information	20, 40, 100
Agents	types of agents	humans, humans + database, humans + avatars

Number of Worlds: 27

[†] When there is a database the total number of agents is the number of humans plus one. When there are avatars the total number of agents is 2 times the number of humans.

Each of the performance metrics could be measured across all agents, rather than across just the human agents. Since not all agents can learn (e.g., avatars) these metrics would not be comparable across the various organizations examined. Therefor for this study these metrics are calculated only across the human agents.

Organizational Communication in the Intelligent Space

As previously noted the basic effect of working in an intelligent space is one of scale — more agents, more information. The second effect is that the availability of information and agents is altered by the presence of artificial agents. How do these changes impact performance?

Simulation results suggest access to more people (group size increases) will result in it taking longer for new information to reach all members of the group and for the group to achieve high accuracy (see figure 1). However, an increase in the population will have little impact on agreement. For information diffusion, there are decreasing costs to scale in population. Thus, in extremely large groups, the main impact of further increases in size will be to decrease the accuracy of decisions. Simulation suggests that access to

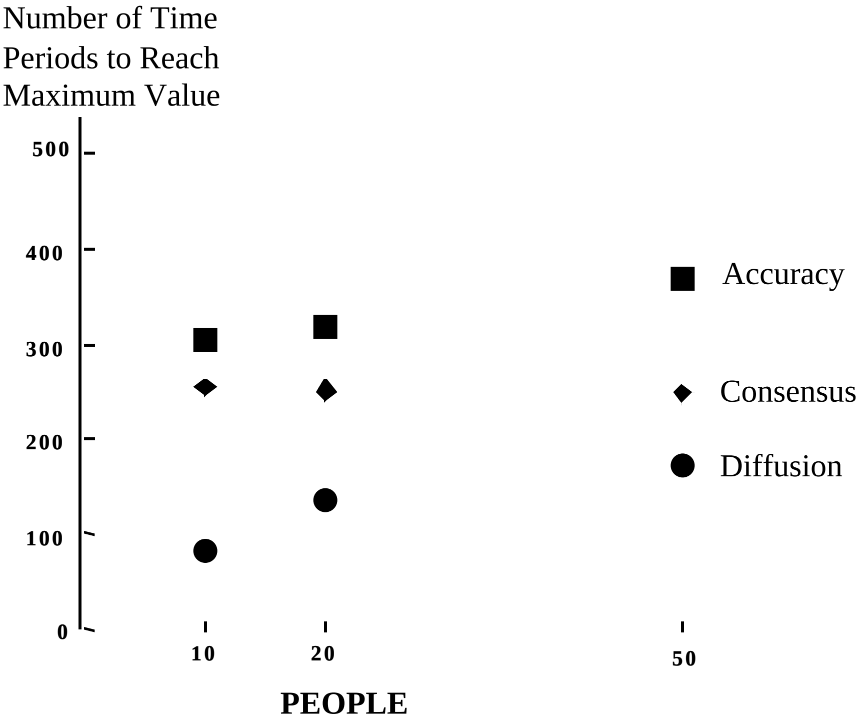

Figure 1. Impact of increase in access to others organizational performance.

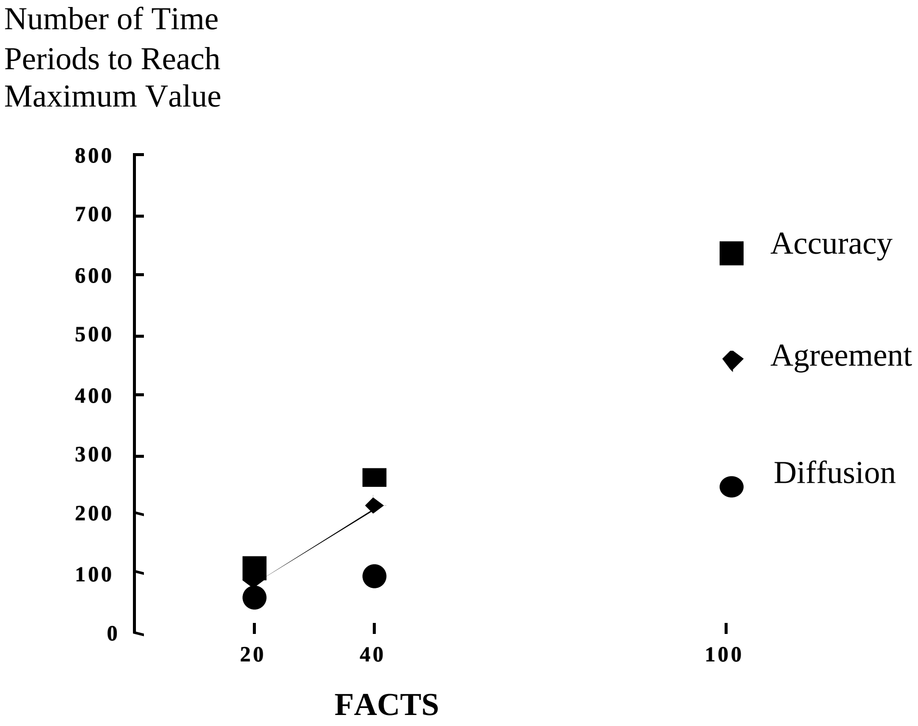

Figure 2. Impact of information explosion on organizational performance.

more information will result in substantial increases in the time it takes information to diffuse, consensus to be reached, and high performance to be achieved (see figure 2). The impact of information growth is much more detrimental than is the impact of population growth on the performance characteristics. These are general first order effects due to the increase in the number of agents and the amount of information to which agents have access. This can be seen even in studies in which there are no smart agents (Carley, 1990; Kaufer and Carley, 1993).

It is typically assumed that in the digital economy everything will occur faster. It should be quicker for individuals to learn new ideas from databases or avatars than to find the person who knows the novel piece of information. Because everyone can access the same information, e.g., in the database, consensus should occur more rapidly and accuracy improved. In contrast, the presence of artificial agents in these intelligent spaces also serves to slow things down (see figure 3). In intelligent spaces, there are now more agents for humans to interact with (databases or avatars). These interactions have little effect on the rate at which information diffuses. However, they do affect the order in which information diffuses; i.e., they affect who learns what when by increasing the number of sources for all information. As a result, they actually make it harder for agreement and high accuracy to be achieved in a timely matter. When artificial agents are not presents, humans act as gatekeepers limiting both the flow of information and who gets what when. This gatekeeping facilitates building agreement in large groups, which can in turn promote higher performance. As anyone can get access to any information any time from anywhere agreement goes down which can reduce accuracy in performance.

Thus far only the direct effects have been examined. There are, however, some interesting interaction effects. Table 4 contains the results of anova analyses looking at the impact of population size, amount of information, and type of agents present (just humans, humans plus database, humans plus avatars) on the performance variables. A word of caution about interpreting these statistical values. In a simulation, N size can be increased arbitrarily. As N size is increased where there is an effect in the limit it becomes significant, and where there is not an effect in the limit it becomes 0. The level of statistical significance is thus used to determine the number of replications to run to achieve a robust result. The value of the coefficients and the R2 is not in whether or not they are significant. Rather, their value is in their relative level. For example, here we see that the factors contributing to diffusion in this non-linear model are simply population size, the amount of information and the size of the knowledge network (population size * amount of information). Of these the amount of information is the dominant factor.

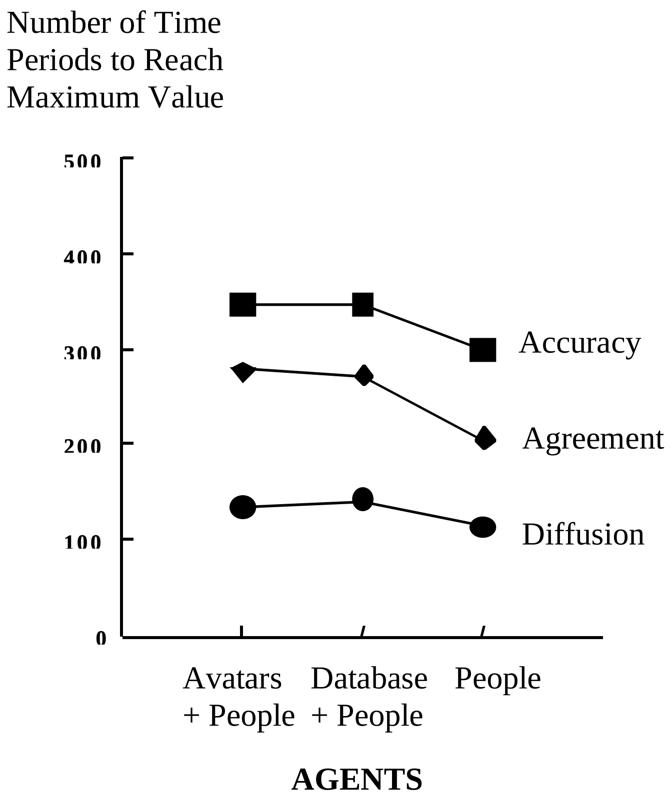

Figure 3. Impact of smart agents on organizational performance.

Table 4. Anova Results

Dependent	Degrees of Freedom	Diffusion	Consensus	Accuracy
Population Size	2	34.90***	0.25	17.20***
Amount of Information	2	168.02***	910.08***	1091.86***
Type of Agents Present	2	3.40	41.15***	11.26***
Population Size * Amount of Information	4	10.691***	0.70	5.69***
Population Size * Type of Agents Present	4	0.08	8.13***	2.60
Amount of Information * Type of Agents Present	4	0.63	16.24***	5.71***
Population Size * Amount of Information * Type of Agents Present	8	1.13	4.38***	1.20
Multiple-R2		0.658	0.893	0.905

Values shown are F-Ratios. *** <= .001, ** = .005, * < .01

N = 270 – 10 runs each of 27 worlds.

For diffusion and accuracy the interaction between facts and population is multiplicative. For consensus, there are interesting interactions with the types of agents and people/facts. Avatars impede agreement when the population is small, whereas databases are more troublesome when the population is large. For large populations, databases and avatars increasingly impede agreement as the fact base grows; but databases are more detrimental than avatars.

Conclusion

Organizations are complex systems. As we move into a world of intelligent spaces artificial smart agents should become increasingly prevalent. Understanding how these agents might alter the form of organizations is imperative. In this study, the basic characteristics of intelligent spaces were described and a combined social network information processing approach to theorizing about the likely effects of such agents was characterized. The complexity of the system was such that computational analysis was used as an aid in thinking through the possible ramifications of these changes. Computational analysis was used to build new concepts, theories, and knowledge about organizations. The computational model is the embodiment of the theory of smart agents and their impact on organizations that being developed and is in need of being tested. Since the model is a model of the underlying process it generate a large number of hypotheses (Carley, 1999c). These hypotheses can then be tested in other settings. Let us consider some of the core hypotheses generated.

The results from the virtual experiments suggest that as we move into intelligent spaces the increase in access to people and information tends to increase the time to achieve high levels of accuracy and for any particular piece of information to diffuse.

Moreover, we saw that, in these simulations, increasing the amount of available information is more devastating on performance than is increasing the number of potential communication partners. It is reasonable to expect that the amount of information people have access to will expand at a greater rate than will the number of people. A variety of factors would contribute to this being the case — increased archiving of information, digitization of old records, education leading to a decrease in population growth, etc. It is often implicitly assumed that access and use are the same; i.e., that if people have access to information they will use it. However, these results suggest that access to more information will actually slow the rate at which any particular piece of information is accessed and learned. Even in the intelligent space, people are still limited in the rate at which they can learn; i.e., bounded rationality still applies. What these results suggest is that increases in competition among ideas, which occurs as the knowledge base expands, can mitigate the value of expanded access.

In these virtual experiments avatars were more helpful in large than in small organizations. An easy explanation is that "sure that's because people would rather talk to people and in a small group you can talk to all the people." However, that is not the causal mechanism. Rather, what is happening in these simulations is that since avatars cannot learn, their value as communicators is higher in a large organization where there are more people from whom to learn. Having agents that do not learn can actually speed communication and overall information diffusion as the same message keeps getting repeated to all interaction partners. Databases become less effective in larger populations as they keep changing and growing in their contents.

These results suggest that for large populations, databases and avatars increasingly impede agreement as the knowledge base grows; but databases are more detrimental than avatars. This is particularly true for accuracy. The type of task used, is essentially a voting task (classification and choice). High accuracy means that all of the personnel correctly identify the code and vote the same. What these results are implying is that for tasks of this sort, such as budget setting, elections, and setting production levels, the movement to intelligent spaces may create a sense of unease. The implications is that if decisions need to be made in the same time span as in a non-digital economy, the level of agreement is likely to be lower and the accuracy of the decision is likely to be lower. In that way, movement to intelligent spaces, may actually lead to a culture of dissent.

If these simple technologies have this affect on organizational performance, how might organizations respond to facilitate rapid diffusion, to build consensus, and to enable the organization to rapidly achieve high accuracy? This analysis provides a few clues. First, the results suggest that agents that cannot learn can speed information diffusion. One way of implementing this in the database world would be to create multiple smaller databases, on special topics, and lock them once they are full rather than allowing people to keep adding to them. Second, these results suggest that a large increase in available information can be disabling. This suggests that when consensus and rapid diffusion of new ideas is the goal, that goal can be facilitated by putting in place procedures to screen information, rate it, or otherwise limit access. Third, one of the reasons that avatars and databases slow things down in these virtual worlds is that they create an effectively more densely connected interaction network. This suggests that factors that promote grouping, gatekeeping, and so forth will actually facilitate information diffusion and consensus. Organizing schemes such as group distribution lists, limited web access, group related pointers between sites, that constrain the interaction network and set boundaries may have the ironic effect of promoting diffusion and the development of consensus within the group.

References

Aldrich, H.E. & S. Mueller, 1982. The evolution of organizational forms: technology, coordination, and control. In *Research in Organizational Behavior*, vol. 4, edited by B.M. Staw & L.L. Cummings. Greenwich, CN: JAI, 33-87.

Barley, Stephan. 1990. "The Alignment of Technology and Structure through Roles and Networks." *Administrative Science Quarterly*, 61-103

Baum, J. A. C. and C. Oliver, 1991, "Institutional Linkages and Organizational Mortality", *Administrative Science Quarterly*, 36: 187-218

Benford, S., Bowers, J., Fahlen, L.E., Greenhalgh, C. & Snowdon, D., 1997. "Embodiments, avatars, clones & agents for multi-user, multi-sensory virtual worlds". *Multimedia Systems*, 5(2): 93-104.

Beniger, J.R. (1986). *The control revolution. Technological and economic origins of the information society*. Cambridge, MA: Harvard University Press.

Blau, P. & W.R. Scott, 1962. Formal Organizations, San Francisco: Chandler.

Bookman, Lawrence A. 1994. *Trajectories through knowledge space : a dynamic framework for machine comprehension*. Boston : Kluwer Academic Publishers.

Burt, Ronald S., 1992. *Structural Holes: the social structure of competition*. Cambridge, MA: Harvard University Press.

Burton, R.M. and B. Obel. 1984. Designing Efficient Organizations: Modeling and Experimentation. Amsterdam: Elsevier Science.

Burton, Richard M. & Borge Obel, 1998. Strategic Organizational Design: Developing Theory for Application, Kluwer Academic Publishers.

Carley, Kathleen M. Michael J. Prietula & Zhiang Lin, 1998, “Design versus Cognition: The Interaction of Agent Cognition and Organizational Design on Organizational Performance, “ Journal of Artificial Societies and Social Simulation, 1(3):1-19. 30 June 1998 at <http://www.soc.surrey.ac.uk/JASSS/>, 1: paper 4.

Carley, Kathleen & Allen Newell, 1994, "The Nature of the Social Agent." *Journal of Mathematical Sociology* , 19(4): 221-262.

Carley, Kathleen & Michael Prietula, 1994, "ACTS Theory: Extending the Model of Bounded Rationality," in Kathleen Carley & Michael Prietula (Eds.), *Computational Organization Theory* .Hillsdale, NJ: Lawrence Earlbaum Associates.

Carley, Kathleen 1992, "Organizational Learning and Personnel Turnover." Organization Science, 3(1): 20-46.

Carley, Kathleen M. & David M. Svoboda, 1996, “Modeling Organizational Adaptation as a Simulated Annealing Process.” *Sociological Methods and Research*, 25(1): 138-168.

Carley, Kathleen M. & Ju-Sung Lee,1998, Dynamic Organizations: Organizational Adaptation in a Changing Environment.Ch. 15 (pp. 269-297) in Joel Baum (Ed.) *Advances in Strategic Management*, Vol 15, Disciplinary Roots of Strategic Management Research. JAI Press. Pp. 269-297.

Carley, Kathleen M. & Les Gasser, 1999, "Computational Organization Theory." Ch. 7 in Gerhard Weiss (Ed.) Distributed Artificial Intelligence. Cambridge, MA: MIT Press.

Carley, Kathleen M. 1999a. Organizational Change and the Digital Economy: A Computational Organization Science Perspective. In Brynjolfsson, Erik and Brian Kahin, Eds., Understanding the Digital Economy: Data, Tools, Research, MIT Press, Cambridge, MA.

Carley, Kathleen M. 1999b. "On the Evolution of Social and Organizational Networks." In Steven B. Andrews and David Knoke (Eds.) Vol. 16 special issue of Research in the Sociology of Organizations. on *"Networks In and Around Organizations*." JAI Press, Inc. Stamford, CT, pp. 3-30.

Carley, Kathleen M. 1999c, "On Generating Hypotheses Using Computer Simulations. " *Systems Engineering*, 2(2): 69-77.

Carley, Kathleen M. 1998, "Organizational Adaptation." *Annals of Operations Research*. 75: 25-47.

Carley, Kathleen M. 1996, "Communicating New Ideas: The Potential Impact of Information and Telecommunication Technology" *Technology in Society*, 18(2): 219-230.

Carley, Kathleen, 1995. "Communication Technologies & Their Effect on Cultural Homogeneity, Consensus, & the Diffusion of New Ideas." *Sociological Perspectives* , 38(4): 547-571.

Carley, Kathleen, 1991. “A Theory of Group Stability.” *American Sociological Review* , 56(3): 331-354.

Carley, Kathleen, 1990, “Group Stability: A Socio-Cognitive Approach.” Pp. 1-44 in Lawler E., Markovsky B., Ridgeway C. & Walker H. (Eds.) *Advances in Group Processes: Theory & Research* . Vol. VII. Greenwhich, CN: JAI Press.

Contractor, N. & Eisenberg, E., 1990. Communication Networks and New Media in Organizations. In *Organizations and Communication Technology*, J. Fulk and C. Steinfield, eds. Newbury Park: Sage Publications.

Cyert, R.M., & J.G. March, 1963, *Behavioral theory of the firm*. Englewood Cliffs, NJ: Prentice-Hall.

Di Martino, V. & L. Wirth, 1990. “Telework: A New Way of Working & Living.” *International Labour Review*, 129(5): 529-554.

Epstein, Josh & Rob Axtell, 1997. *Growing Artificial Societies*, Boston, MA: MIT Press

Etzioni A. 1964. *Modern Organization*, Englewood Cliffs, NJ: Prentice-Hall.

Fligstein, Neil, 1985.The spread of the multi-divisional form among large firms, 1919-1979. *American Sociological Review* , 50: 377-391.

Harris, D.H. (Ed.) 1994,*Organizational Linkages: Understanding the Productivity Paradox*. Washington D.C.: National Academy Press..

Hutchins, Edwin. 1991. "The Social Organization of Distributed Cognition." In *Perspectives on Socially Shared Cognition*, edited by L. B. Resnick, J. M. Levine and S. D. Teasley. Washington D.C.: American Psychological Association.

Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.

Jin, Yan & Raymond Levitt, 1996. The Virtual Design Team: A computational model of project organizations. Computational and Mathematical Organization Theory}, 2(3): 171-196.

Kaufer, David S. & Kathleen M. Carley, 1993, *Communication at a Distance: The Effect of Print on Socio-Cultural Organization & Change*, Hillsdale, NJ: Lawrence Erlbaum Associates.

Kiesler, Sara (Ed.), 1996. *Culture of the Internet.* Mahwah, N.J. : Lawrence Erlbaum Associates.

Krackhardt, David & Kathleen M. Carley, 1998, "A PCANS Model of Structure in Organization" Pp. 113-119 in *Proceedings of the 1998 International Symposium on Command and Control Research and Technology*. June. Monterray, CA.

Krackhardt, David and Martin Kilduff. 1994. “Bringing the Individual Back In: A structural analysis of the internal market for reputation in organizations.” Academy of Management Journal 37(1): 87-108.

Kurzweil, Raymond, 1988. *The Age of Intelligent Machines*. Cambridge, Mass. : MIT Press, 1988.

Lant, Theresa L. 1994. "Computer Simulations of Organizations as Experimental Learning Systems: Implications for Organization Theory." Pp. 195-216 in *Computational Organization Theory,* edited by K.M. Carley and M.J. Prietula. Hillsdale, NJ: LEA.

Levinthal, Daniel A. & James G. March. 1981. A model of adaptive organizational search, *Journal of Economic Behavior and Organization*, 2, 307-333.

Levitt, Raymond E., Geoff P. Cohen, John C. Kunz, Chris I. Nass, Torre Christiansen & Yan Jin, 1994. "A Theoretical Evaluation Of Measures Of Organizational Design: Interrelationship And Performance Predictability." Pp. 1-18 in *Computational Organization Theory*, edited by K.M. Carley and M.J. Prietula . Hillsdale, NJ: Lawrence Erlbaum Associates.

Lipnack, Jessica & Jeffrey Stamps, 1997. *Virtual teams : reaching across space, time, and organizations with technology* . New York : John Wiley & Sons.

March, James G. & H. Simon, 1958, *Organizations*. New York: John Wiley and Sons, Inc.

March, James G. 1996. "Exploration and Exploitation in Organizational Learning." In M.D. Cohen & L.S. Sproull (Eds.) Organizational Learning. Thousand Oaks, CA: Sage

March, James G. 1996. "Exploration and Exploitation in Organizational Learning." In M.D. Cohen and L.S. Sproull (Eds.) *Organizational Learning*. Thousand Oaks, CA: Sage

Masuch, M. and P. LaPotin, 1989. Beyond Garbage Cans: An AI Model of Organizational Choice. *Administrative Science Quarterly*. 34: 38—67.

McPherson, J. Miller and Lynn Smith-Lovin. 1987. "Homophily in Voluntary Organizations: Status Distance and the Composition of Face-to-Face Groups." *American Sociological Review* 52: 370-79.

Miles, R. E., & C.C. Snow, 1995. The new network firm: A spherical structure built on a human investment philosophy. *Organizational Dynamics, 23,* 5-18.

Mintzberg, H. (1983). Structures in Five: Designing Effective Organizations. Prentice Hall Inc.

Monge, P.R., & Contractor, N. (in press). Emergence of communication networks. In F.M. Jablin & L.L. Putnam (Eds.), *New handbook of organizational communication*. Thousand Oaks, CA: Sage Publications.

Moreland, R. L., Argote, L., & Krishnan, R. (1996), Socially shared cognition at work: Transactive memory and group performance, In J. L. Nye & A. M. Brower (Eds.), *What's social about social cognition? Research on socially shared cognition in small groups* (pp. 57-84), Thousand Oaks, CA: Sage

Newell, A., 1990. *Unified Theories of Cognition.* Cambridge, MA: Harvard University Press.

Nixon, Paddy, Gerard Lacey, and Simon Dobson (eds.), 1999. Managing Interactions in Smart Environments: 1st International Workshop on Managing Interactions in Smart Environments (MANSE'99), Dublin, Ireland, December 1999 .

Nohria, Nitin & Robert G. Eccles (Eds.), 1992. *Networks and organizations : structure, form, and action*, Boston, MA : Harvard Business School Press.

Powell, W.W. 1990. “Neither Markets nor Hierarchy: Network Forms of Organizations”, in *Research in Organizational Behavior*, 12, 295-336, B.M. Staw and L. L. Cummings (Eds.), JAI Press, Greenwich, CT.

Samuelson, D. 2000. Designing Organizations. OR/MS Today. December.: 1-4. See also http://www.lionhrtpub.com/orms/orms-12-00/samuelson.html

Simon, H., 1955, A behavioral model of rational choice. *Quarterly Journal of Economics,* 69: 99-118.

Simon, H., 1956. Rational choice & the structure of the environment. *Psychological Review,* 63: 129-138.

Sproull, L. and S. Kiesler, 1991. *Connections: New Ways of Working in the Networked Organization*. Cambridge, MA: MIT Press.

Thomas, Peter and Hans-W. Gellersen (eds.), 2000. Proceedings of the International Symposium on Handheld and Ubiquitous Computing: Second International Symposium, HUC 2000, Bristol, UK, September 25-27, 2000 .

Thrun, S., 1996. To Know or Not To Know: The Role of Models in Mobile Robotics. *AI Magazine* 18:1.

Tushman, M.L. & Nadler, D. , 1978, Information processing as an integrating concept in organizational design. *Academy of Management Review*, 3: 613-624.

Weber, Max,1947, *The Theory of Social & Economic Organization*, A. H. Henderson & Talcott Parsons (eds.). Glencoe, IL: Free Press, pp. 329-341.

Weiss, Gerhard (Ed.), 1999. *Distributed Artificial Intelligence*. Cambridge, MA: MIT Press.

White, H., 1970. *Chains of Opportunity*. Cambridge, Mass: Harvard.

Worthington, John, 1997. *Reinventing the workplace*. Oxford ; Boston : Architectural Press, 1997.

Yamagata, H., K. S. Yeh, S. Stewman & H. Dodge. 1997. "Sex Segregation and Glass Ceilings: A Comparative Statics Model of Women's Career Opportunities in the Federal Government over a Quarter of a Century." American Journal of Sociology 103: 566-632.

---

[i] This is a reduced version of the PCANS formulation first proposed by Krackhardt and Carley, 1998. They defined a meta-matrix of relations among people, resources and tasks. Without loss of generality we redefine resources as knowledge and expand people to include all intelligent agents.