ENHANCED KNOWLEDGE-BASED HUMAN ORGANIZATION AND SOCIAL CHANGE

Kathleen M. Carley, Carnegie Mellon University

Changes that bring together nanotechnology, information science, biology, and cognition have the potential to revolutionize the way we work and organize society. A large number of outcomes are possible. At the same time, existing social forms, legislation, and culture will limit and direct the potential outcomes. In a very real sense, technologies and societies, tools and cultures, capabilities and legislation will co-evolve. Without attempting to predict the future, a series of possible outcomes, issues, and research challenges are discussed. Particular emphasis is placed on issues of security and potentially radical change within groups, organizations, and society.

Data and Privacy

In the area of bioterrorism, a key issue is early detection or "biosurveillance." Early detection requires smart sensors at the biological level in the air, water, and ground and on humans. Early detection requires integrating this data with geographic, demographic, and social information. Even were the sensors to exist, there would still be a problem: Under current legislation and privacy laws, the data cannot be integrated and made readily accessible to practitioners and researchers. To develop and test data mining tools, knowledge management tools, and what-if policy simulators, access is needed to a wide range of data in real time; but, providing access to such data enables the users of these tools to "know" details of individual behavior.

In the area of organizations, a key issue is team design and redesign (Samuelson 2000). Team design and redesign requires accurate data of who knows what, can work with whom, and is currently doing what. Doing such a skill audit, network analysis, and task audit is a daunting task. Maintaining the information is even more daunting. Individuals are loathe to provide the information for fear of losing their basis of power or anonymity, or for fear of reprisal. However, much of the information is implicit in the locations that people occupy, their stress levels, webpages, curricula vitae, public conversations, and so on.

In the cases of both acquiring and maintaining individual data, all of the following can be used to enable better outcomes: nano-bio-sensors that are embedded in the body and that report on individual health, stress level, and location; intelligent surfaces that track who is present while reshaping themselves to meet the needs of and enhance the comfort of the users; auto-sensors that create a memory of what is said when people cough or sneeze; air and water sensors that sense contaminants; data-mining tools that locate information, simulation tools that estimate the change in social outcomes; information assurance tools and secure distributed databases. Indeed, such tools are critical to the collection, analysis, protection, and use of information to enhance group performance. The relatively easy problems here will be those that are dominated by technology, e.g., distributed database tools, data integration procedures, information assurance technology, and smart sensors. Those problems dealing with the need to change cultures, legislation, and ways of working will be more difficult. Privacy laws, for example, could mitigate the effectiveness of these tools or even determine whether they are ever developed. There are many critical privacy issues, many of which are well identified in the NRC report, *The Digital Dilemma* (http://www.nap.edu/catalog/9601.html). Views of knowledge as power will limit and impede data collection. Having such data will revolutionize healthcare, human resources, career services, intelligence services, and law enforcement. Having such data will enable "big-brotherism." Were we able to overcome these two mitigating factors, then a key issue will become, "What will the bases for power be when knowledge is no longer a controlled commodity?" Since many organizations are coordinated and managed through the coordination and management of information, as knowledge is no longer controlled, new organizational forms should emerge. For example, a possible result might be the development of monolith corporations with cells of individuals who can do tasks, and as those tasks move from corporation to corporation, the cells would move as well. In this case, benefits, pay scales, etc., would be set outside the bounds of a traditional corporation. In this case, individual loyalty would be to the area of expertise, the profession, and not the company. Corporations would become clearinghouses linking agents to problems as new clients come with new problems.

Ubiquitous Computing and Knowledge Access

As computers are embedded in all devices, from pens to microwaves to walls, the spaces around us will become intelligent (Nixon, Lacey, and Dobson 1999; Thomas and Gellersen 2000). Intelligent spaces are generally characterized by the potential for ubiquitous access to information, people, and artificial agents, and the provision of information among potentially unbounded networks of agents (Kurzweil 1988). The general claim is that ubiquitous computing will enable everyone to have access to all information all the time. In such an environment, it is assumed that inequities will decrease. This is unlikely. While ubiquitous computing will enable more people to access more information more of the time, there will still be, short of major reforms, people with little to no access to computing. There will be excess information available, information making it difficult to discern true from false information. There will be barriers in access to information based on legislation, learning, and organizational boundaries. While information will diffuse faster, the likelihood of consensus being reached and being accurate given the information will depend on a variety of other factors such as group size, the complexity of the task and associated knowledge, initial distribution of information in the group, and so on. As a result, things may move faster, but not necessarily better.

Initial simulation results suggest that even when there are advanced IT capabilities, there will still be pockets of ignorance, certain classes of individuals will have privileged access to information and the benefits and power that derive from that, groups will need to share less information to be as or more effective, databases may decrease shared knowledge and guarantee information loss, and smaller groups will be able to perform as well or better than larger groups (Alstyne, M. v., and Brynjolfsson, E. 1996; Carley 1999). To address issues such as these, researchers are beginning to use multiagent network models. These models draw on research on social and organizational networks (Nohira and Eccles 1992), advances in network methodology (Wasserman and Faust 1994), and complex system models such as multiagent systems (Lomi and Larsen 2001). In these models, the agents are constrained and enabled by their position in the social, organizational, and knowledge networks. These networks influence who interacts with whom. As the agents interact, they learn, which in turn changes with whom they interact. The underlying networks are thus dynamic. The results suggest that organizations of the future might be flatter, with individuals coming and going from teams based on skills, that is, what they know, and not whom they know. As a result, social life will become more divorced from organizational life. Initial simulation results suggest that if information moves fast enough, decisions will become based not as much on information as on the beliefs of others; this should be particularly true of strategic decisions.

Socially Intelligent Technology

Major improvements in the ability of artificial agents to deal with humans and to emulate humans will require those artifacts to be socially intelligent. Socially intelligent agents could serve as intelligent tutors, nannies, personal shoppers, etc. Sets of socially intelligent agents could be used to emulate human groups/organizations to determine the relative efficacy, feasibility, and impact of new technologies, legislation, change in policies, or organizational strategy. At issue are questions of how social these agents need to be and what is the basis for social intelligence. It is relatively easy to create artificial agents that are more capable than a human for a specific well-understood task. It is relatively easy to create artificial agents that can, in a limited domain, act like humans. But these factors do not make the agents generally socially intelligent. One of the research challenges will be for computer scientists and social scientists to work together to develop artificial social agents. Such agents should be social at both the cognitive and precognitive (bio) level. Current approaches here are software- limited. They are also potentially limited by data; nanotechnology, which will enable higher levels of storage and processing, will also be necessary. That is, creating large numbers of cognitively and socially realistic agents is technically unfeasible using a single current machine. Yet, such agents need to exist on a single machine if we are to use such tools to help individuals manage change.

A key component of social intelligence is the ability to operate in a multiagent environment (Epstein and Axtell 1997; Weiss 1999). However, not all multiagent systems are composed of socially intelligent agents. For a machine to be socially intelligent, it needs to be able to have a “mental” model of others, a rich and detailed knowledge of realtime interaction, goals, history, and culture (Carley and Newell 1994). Socially intelligent agents need transactive memory, i.e., knowledge of who knows whom (the social network), who knows what (the knowledge network), and who is doing what (the assignment network). Of course this memory need not be accurate. For agents, part of the “socialness” also comes from being limited cognitively. That is, omniscient agents have no need to be social, whereas, as agents become limited — boundedly rational, emotional, and with a specific cognitive architecture — they become more social.

One of the key challenges in designing machines that could have such capabilities is determining whether such machines are more or less effective if they make errors like humans do. What aspects of the constraints on human cognition, such as the way humans respond to interrupts, the impact of emotions on performance, and so on, are critical to acquiring and acting on social knowledge? While we often see constraints on human cognition as limitations, it may be that social intelligence itself derives from these limitations and that such social intelligence has coordinative and knowledge benefits that transcend the limitations. In this case, apparent limits in individuals could actually lead to a group being more effective than it would be if it were composed of more perfect individual agents (Carley and Newell 1994).

A second key challenge is rapid development. Computational architectures are needed that support the rapid development of societies of socially intelligent agents. Current multiagent platforms are not sufficient, as they often assume large numbers of cognitively simple agents operating in a physical grid space as opposed to complex intelligent, adaptive, learning agents with vast quantities of social knowledge operating in social networks, organizations, and social space. Moreover, such platforms need to be extended to enable the co-evolution of social intelligence at the individual, group, and organizational level at differing rates and to account for standard human processes such as birth, death, turnover, and migration.

A third challenge is integrating such systems, possibly in real time, with the vast quantities of data available for validating and calibrating these models. For example, how can cities of socially intelligent agents be created that are demographically accurate, given census data?

Socially Engineered Intelligent Computer Anti-Viruses and DDOS Defenses

Computer viruses have caused significant financial losses to organizations (CSI 2000). Even though most organizations have installed anti-virus software in their computers, a majority of them still experience infections (ICSA 2000). Most antivirus software can not detect a new virus unless it is patched with a new virus definition file. New virus countermeasures have to be disseminated once a new virus is discovered. Studies of viruses demonstrate that the network topology and the site of the initial infection are critical in determining the impact of the virus (Kephart 1994; Wang 2000; Pastor-Satorras 2001). What is needed is a new approach to this problem. Such an approach may be made possible through the use of socially intelligent autonomous agents.

The Web and the router backbone can be thought of as an ecological system. In this system, viruses prey on the unsuspecting, and distributed denial of service attacks (DDOS) spread through the networks "eating" or "maiming" their prey. Viruses are, in a sense, a form of artificial life (Spafford 1994). One approach to these attacks is to propagate another "species" that can in turn attack these attackers or determine where to place defenses. Consider a computer anti-virus. Computer anti-viruses should spread fixes and safety nets, be able to "eat" the bad viruses and restore the machines and data to various computers without, necessarily, the user's knowledge. Such anti-viruses would be more effective if they were intelligent and able to adapt as the viruses they were combating adapted. Such anti-viruses would be still more effective if they were socially intelligent and used knowledge about how people and organizations use computers and who talks to whom in order to assess which sites to infiltrate when. We can think of such anti-viruses as autonomous agents that are benign in intent and socially intelligent.

Social Engineering

Combined nano-, bio-, info-, and cogno-technologies make it possible to collect, maintain, and analyze larger quantities of data. This will make it possible to socially engineer teams and groups to meet the demands of new tasks, missions, etc. The issue is not that we will be able to pick the right combination of people to do a task; rather, it is that we will be able to pick the right combination of humans, webbots, robots, and other intelligent agents, the right coordination scheme and authority scheme, the right task assignment, and so on, to do the task while meeting particular goals such as communication silence or helping personnel stay active and engaged. Social engineering is, of course, broader than just teams and organizations. One can imagine these new technologies enabling better online dating services, 24/7 town halls, and digital classrooms tailored to each student's educational and social developmental level.

The new combined technologies are making possible new environments such as smart planes, "living" space stations, and so on. How will work, education, and play be organized in these new environments? The organizational forms of today are not adequate. Computational organization theory has shown that how groups are organized to achieve high performance depends on the tasks, the resources, the IT, and the types of agents. You simply do not coordinate a group of humans in a board room in the same way that you would coordinate a group of humans and robots in a living space station or a group of humans who can have embedded devices to enhance their memory or vision.

Conclusion

These areas are not the only areas of promise made possible by combining nano-, bio-, info-, and cogno-technologies. To make these and other areas of promise turn into areas of advancement, more interdisciplinary research and training is needed. In particular, for the areas listed here, joint training is needed in computer science, organizational science, and social networks.

References

Alstyne, M.v., and E. Brynjolfsson. 1996. Wider access and narrower focus: Could the Internet Balkanize science? *Science* 274(5292):1479-1480.

Carley, K.M. forthcoming, Smart agents and organizations of the future. In *The handbook of new media*, ed. L. Lievrouw and S. Livingstone.
_____. forthcoming. Computational organization science: A new frontier. In Proceedings, Arthur M. Sackler Colloquium Series on Adaptive Agents, Intelligence and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling, October 4-6, 2001; Irvine, CA: National Academy of Sciences Press.
_____. forthcoming, Intra-Organizational Computation and Complexity. In *Companion to Organizations,* ed. J.A.C. Baum. Blackwell Publishers.
Carley, K.M., and V. Hill. 2001. Structural change and learning within organizations. In *Dynamics of organizations: Computational modeling and organizational theories,* ed. A. Lomi and E.R. Larsen. MIT Press/AAAI Press/Live Oak.
Carley, K.M. 1999. Organizational change and the digital economy: A computational organization science perspective. In *Understanding the Digital Economy: Data, Tools, Research*, ed. E. Brynjolfsson, and B. Kahin. Cambridge, MA: MIT Press.
Carley, K.M., and A. Newell. 1994. The nature of the social agent. *J. of Mathematical Sociology* 19(4): 221-262.
CSI. 2000. CSI/FBI computer crime and security survey. *Computer Security Issues and Trends.*
Epstein, J., and R. Axtell. 1997. *Growing artificial societies.* Boston, MA: MIT Press.
ICSA. 2000. ICSA Labs 6th Annual Computer Virus Prevalence Survey 2000. ICSA.net.
Kephart, J.O. 1994. How topology affects population dynamics. In *Artificial life III,* ed. *C.G.* Langton. Reading, MA: Addison-Wesley.
Kurzweil, R. 1988. *The age of intelligent machines*. Cambridge, MA: MIT Press.
Lomi, A., and E.R. Larsen, eds. 2001. *Dynamics of organizations: Computational modeling and organizational theories*. MIT Press/AAAI Press/Live Oak.
Nixon, P., G. Lacey, and S. Dobson, eds. 1999. Managing interactions in smart environments. In Proceedings, 1st International Workshop on Managing Interactions in Smart Environments (MANSE '99), Dublin, Ireland, December 1999.
Nohira, N. and R. Eccles, eds. 1992. *Organizations and networks: Theory and practice*. Cambridge, MA: Harvard Business School Press.
Pastor-Satorras, R., and A. Vespignani. 2001. Epidemic dynamics and endemic states in complex networks. Barcelona, Spain: Universitat Politecnica de Catalunya.
Samuelson, D. 2000. Designing organizations. *OR/MS Today*. December: 1-4. See also http://www.lionhrtpub.com/orms/orms-12-00/samuelson.html.
Spafford, E.H. 1994. Computer viruses as artificial life. *Journal of Artificial Life*.
Thomas, P. and H.-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 .
Wang, C., J.C. Knight, and M.C. Elder. 2000. On computer viral infection and the effect of immunization. In Proceedings, IEEE 16th Annual Computer Security Applications Conference.
Wasserman, S. and K. Faust. 1994 *Social Network Analysis.* New York: Cambridge University.
Weiss, G., ed. 1999. *Distributed artificial intelligence*. Cambridge, MA: MIT Press.