The Model Class Discovery Dilemma in Computer-Supported Work Environments —From Critical Incidents to Metrics of Coordination

by Rodney Fuller
Department of Psychology
Claremont Graduate School
Claremont, Ca 91711 USA
<fullerr@cgsvax.claremont.edu>

Abstract

Given limited resources the main problem in designing tools for computer-supported work environments is to quickly find both a level of representation and a level of analysis that effectively explains the uncertainty that exists within the resulting "electronic space". This is also known as "the model class discovery dilemma": the problem of finding a fit between design and environmental variability. The traditional design paradigm for the field of computer-supported cooperative work (CSCW) has proposed that modelers use either critical incidents or some sort of design rationale to guide the users of these systems toward productive synchronous and asynchronous work behaviors. A critique of the critical incident and design schema methods is presented based on a case study of my designing a distributed CSCW system in the customer service department of a large health maintenance organization. A third alternative that looks at designing based on more fundamental "metrics of coordination" is also presented. The implications that computational systems utilizing new coordination metrics might have on the sociotechnical work structure of the future is also presented.

Group work environments are constrained by time, motivation, and the availability of and strategies related to the use of internal and external information processing resources. If the group structure remains relatively stable over time then there will be the added complexity of within group variations in expertise and commitment. This paper looks at the problems related to creating computer models and tools that try and support the above complex social, cognitive and technical environment.

Given that groups of agents will always be constrained by both internal and environmental factors designers need to find ways of constructing models of group work that do not result in the breakdown of workflow or communication between the actors. This reality requires that modelers deal with what has become known as the "representation problem" within the field of Artificial Intelligence and the "level of analysis" problem within Organizational Behavior. Both of these problems are seen as fundamental tradeoffs with their respective fields, but within the field of Statistical Mechanics these problems are combined into the "Model Class Discovery Dilemma" and several solutions exist for resolving this dilemma: the model class discovery dilemma is that dilemma that all designers face of determining when they have represented enough of the variability in the environment so that a formal design of that understanding will result in a useful system. A system that tries to do too much is slow and cumbersome, while a system that does not do enough is equally worthless because actors are not able to obtain, manipulate and produce solutions that reduce the variability within the model. Jim Crutchfield examines this dilemma from a more theoretical perspective:

> "The epistemological problem of nonlinear modeling is: Have we discovered something on our data or have we projected the new found structure onto it? This was the main lesson of attempting to reconstruct equations of motion from a time series: When it works, it works; When it doesn't, you don't know what to do; and in both cases it is ambiguous what you have learned. Even though data was generated by well-behaved, smooth dynamical systems, there was an extreme sensitivity to the assumed model class that completely swamped 'model order estimation.' Worse still there was no *a priori* way to select the class appropriate to the process. This should be contrasted with what is probably one of the more important practical results in statistical modeling: within a model class a procedure exists to find, given a finite amount of data, an optimal model that balances prediction error against model complexity. Despite representations to the contrary, this 'model order estimation' does not address issues of class inappropriateness and what to do when confronted with failure" (Crutchfield 1992a, p. 68).

Traditional designers of Computer-Supported Cooperative Work (CSCW) systems have recommended the use of both critical incidents and design rationales for constructing models of user behavior. Both of these recommendations fail to consider the complexity of the work environments that they hope to model and thus are likely to be only partially successful. Using examples taken from my own experiences designing a small groupware system for the customer service department of a large health maintenance organization (HMO) this paper will discuss the problems related to the "model class discovery dilemma" as they relate to building distributed CSCW systems utilizing the following design representations: 1) critical incidents as proposed by sociotechnical design and participatory design, 2) design scenarios, and 3) present a research proposal for designing "metrics of coordination" that allow for a more adaptive approach to the model class discovery dilemma. Finally the sociotechnical implications that metrics of coordination have in regard to the future of automated work environments will be discussed.

The Case: A CSCW system in a large Health Maintenance Organization:

I recently completed leading a design team in building a simple groupware system that would support the customer service department of a very large health maintenance organization in Southern California. The main role of the system was to support the front-line employees in documenting and researching complaints from customers and possible legal issues related to poor health care or service. A second goal for the system was to vastly improve the quality and timeliness of the workload statistics that the managers of these front line employees had to work with. These two system "goals" were incompatible because the nature of the work that needed to be supported was mainly cooperative in nature, while the nature of the goal that supported the managerial system was procedural (information was to be used to schedule employees and insure compliance with changing governmental criteria).

The basic system allowed for front line employees to log complaints and track performance and collect information related to that complaint. There were 162 employees that used this system and about 100,000 complaints were logged into this system every year. The complaints varied widely (from "I can't find a place to park" to "you took out the wrong kidney") and resources had to be managed dynamically in each of the 13 local medical centers with customer service departments. The ability to distribute work based on the workload and experience of each employee was as a major motivation to use the system in those offices where the most experienced employees were sent out to satellite offices to work alone (they could handle more variety, but they were too valuable to be left idle because of lack of work).

A generic complaint might enter the system in the following way: At any time someone might call (or walk into) the customer service department and complain. Often several separate issues were delivered in one "complaint." To better predict the workload of each office the managers had classified these issues into categories of severity (potential cost to company and potential cost in employees time). In order to support the front line worker these connected issues had to be bundled together to give context to the complaint as a whole (much as they would be presented in a court of law). To support the managers each issue had to be processed separately. Attempts to determine how the issues should be processed for each group of "users" resulted in too many solutions to implement effectively. The main problem was one of perspective—to do their job best the front-line staff had to be left alone to coordinate with the other departments within the organization—while the managers needed hourly updates regarding the workload (both short term and long term) of their employees to make them more productive.

In designing the prototypes and evaluating the usablity of the system I used two different recommendations that other researchers in the field of designing effective CSCW have used: the use of critical incidents and the use of design rationales. The advantages and disadvantages that each of these paradigms presented are summarized below:

Using Critical Incidents as Metrics of Design:

In Computer Supported Cooperative Work (CSCW) there must be a means for justifying the rationale of any work involving distributed decision making—especially decision making that involves several levels of process control. One way that designers can capture these rationales is by using the participatory design techniques described by Greenbaum (1993): "these techniques stress the need for system developers to learn from the experiences of people using computers, not just from formal system descriptions of work. In addition, cooperative approaches argue that workplace language and daily experience of users need to be placed center stage in an effort to enable users" (p. 31).

Designers who rely on user feedback (Norman and Draper, 1986) or participatory design (Greenbaum and Kyng, 1991) to structure or justify rationales are likely to find that the usability of their design will be limited by a failure to include considerations of coordination costs between people, technology and process controls. Techniques for knowledge elicitation based on ethnographic methods will most likely extract examples of behaviors and rationales that are highly situational. Designers who rely on only this type of knowledge when structuring their design environment will build systems where the flow of behavior from one incident to another does not occur without breakdown.

In designing the groupware system I went to three out of the 13 offices and held a three day participatory design session. Most of the first day was taken up in defining the basic nature of the task and constructing some paper prototypes of how the system should work. The basic structure of the overview meeting followed Spradley's (1979) grand view model of ethnographic questioning. In these meetings each office was asked to generally describe how they collect and process complaints received by customers. Once an overall structure had been elicited I then asked more and more specific questions regarding the behaviors that needed to be performed. Once most of the behaviors had been identified I asked the participants to draw a picture of how they wanted their screens to look and act. These drawings were converted into a simple prototype for the second day where further domain elicitation and the start of user testing occurred. Based on the information obtained from the second day further functionality was built into the system and all those people who had not participated in the design process were trained on using the system. In ever office that did not participate in the design process users were encouraged to make suggestions on how they might have structured the design differently. In many cases I was able to make the changes while they waited (changes in the order in which different fields were selected, or how information appeared in those fields) but deeper recommendations regarding the structure of information captured and the nature of process flow were documented for further versions.

The second prototype resulted from the need to combine the 13 different systems into one supportable system. In doing this I first identified all those areas where the individual office system were different and questioned the offices that differed regarding the rationales they used for those processes. The resulting system was installed and accepted (i.e. used) in all of the offices about 7 months after the first prototype was installed. In resolving the differences between the different systems I noticed that the quality of information I got regarding rationales from each office varied widely. Specifically the rationales obtained from offices that had a more autocratic structure tended to focus on incidents that were uncommon and high profile while the more democratic offices provided a more balanced view of the work environment. This difference in perception (and hence elicitation of design space) was caused by the different ways that each office reviewed the productivity of their employees—autocratic offices tended to emphasize those events that were outstanding with praise and raises, while the democratic offices were perceived as rewarding consistent performance.

In summary the elicitation of critical incidents using ethnographic and participatory design paradigms allowed me to quickly prototype a system for documenting work. But in addition to the many benefits that these techniques bring there is also the danger that other factors within the office (but still outside the participatory analysis) can influence the quality of the information obtained. The elicitation of critical incidents is most likely to build a domain model where the only possibilities considered are those that are currently being used. The integrity of this design space will be influenced by any factor that could artificially increase the distance between these incidents. It would be much more beneficial to design for situational use rather than try and support the clusters of elicited (and somewhat unconnected) behaviors. This problem is also present when modelers of business processes (e.g. business process re-engineering) try and model the future possibilities of an organizational design based on an understanding of present practices. The use of design scenarios is one example of how one can implement the use of scripts and plans into designs that support cooperation at different levels.

Design Rationales as Streams or Networks of Incidents:

One of the problems with the use of critical incidents is that the incidents become islands of representation within the design space of all possible work procedures. The incident is not balanced by measures of likelihood, and depending on whether or not the design team is interdisciplinary the information extracted may or may not lack perspective. The inability to use knowledge representations to support multiple layers of interaction can become a serious problem in systems that support complex work environments. Such was the case with my system. In the later prototypes I started using previous histories of incidents to determine if the cost of not including that incident would significantly affect the flow of work. Using this simple method I was able to determine if there should not be a complaint code for something that had happened only once seven years ago, regardless of the fact that the issue received local media attention. But this problem became more complex when other informational tools were added into the final system.

The mainframe system included several other tools that had not been available on-line in the prototypes. These tools included access to the clinical appointment database, the billing and accounting databases, the database regarding special benefits and programs, the database regarding the past history of complaints and access to the complaints that the other offices were currently working on. These additional tools made both the management of on-line tools and the management of employees workflow much more difficult. In order to give the managers a way of tracking the workload of each employee it became necessary to give every employee the capability to immediately list and manage their unfinished work. The ability to pull up a list of all of their outstanding cases was most successful in those offices that had previously required employees to manage their own work, but the added capability of letting the employees prioritize their outstanding work increased the usability of the system based on employee's self-reports, analysis of workflow and measures of general productivity.

The more complex problem was that the productivity of the system seemed to fluctuate depending on highly situational factors. There seemed to be a critical band of inputs (the number of service issues logged) between which the system was productive, but beyond which the productivity dipped below the company standard. One possibility that was investigated was that when very few complaints were being logged into the system that the employees were using the system to train themselves by exploring the full functionality of the system. For example: they might investigate every unpaid bill for someone who was complaining about there not being enough parking places. This training in times of few complaints would lead a employee to try and use all of the tools available to them when they had a full queue of customers waiting to complain. Interviews with and observation of employees showed that users were not returning to check previous work. Two alternatives for dealing with this problem were proposed: 1) that either the employees were not able to remember to go back and resolve complaints after their current screen had disappeared, or 2) that in times of high use some of the more esoteric functionality of the system should be made "temporarily unavailable." Both of these problems required the construction of complex models of behaviors that would link critical incidents together—something which is also called using design scenarios.

The problem I found related to constructing design scenarios for strings of complaints that would either A) automatically set a priority on any unfinished issue and then "tickle" the user to go back and finish their work, or B) to limit the functionality of the system in times of high use, was that the situational applicability of the scenario was not stable. Even though the range of possible issues was clear, the number of ways that these events could combine into a complaint and the number of organizational resources that needed to be contacted were far beyond the power of the system to model. It was thought that a more general theory of coordination needed to be developed. Tom Malone (1988) proposes that a small set of semantic business processes be modeled and that CSCW systems take advantage of the coordination costs between these processes when guiding users and designers. Design Scenarios seem particularly ill-suited to this because they offer a non-adaptive structure from which to predict interaction. Callan et. al. (1991a and 1991b) and others (Sycara 1993; Holsapple, Pakath, Jacob and Zaveri 1993) have shown that adaptive architectures that use machine learning algorithms to set the number of levels of analysis looked at are more productive than more traditional based methods of analysis. In the next section I will discuss what it might mean to take this idea even further into a theory based on "metrics of coordination".

Constructing "Metrics of Coordination"

What is needed is a model of design that allows for the following assumption: That there is no *a priori* means to determine with absolute certainty which factor[s] will explain the most variance surrounding any given event or set of events. As discussed above there are two types of general problems why this is true—a problem related to the representation of information and a problem related to the analysis of multiple levels of meaning. These two problems are discussed below:

The problem of scale is well known in AI as the representation problem, but this problem is not as widely recognized in Organizational Behavior. This problem can be viewed as that decision making process which selects the criteria which will be used later to define the final product[s]. For example in the field of artificial intelligence (AI) it is important to engineer into the data representations slots that will hold all of the information that could possibly influence the computed products, and in organizations it is important to determine, not only what aspects of the available resources (material, capital, information and human) should be managed, but also what combination of criteria should be used as business process indicators (to measure quality, productivity, reliability, etc.).

The second problem is that of scope. While AI researchers have long understood that the selection of a scale limits the scope of the resulting model, many AI researchers seem to have ignored the fact that the selection of such a scale also adds an indeterminate amount of uncertainty into the process as well as limits the ultimate complexity of the system. This issue is well known in Organizational Behavior as "the level of analysis problem" where the basic premise is that the fundamental nature of the problem changes depending upon the number and nature of the participants or perspectives involved in the process (Katz & Kahn 1978, Morgan 1986, Weick 1979). In essence the problem becomes: If one is able to determine a scale from which to operate then this scale imposes both upper and lower limits upon the type and quality of processes that can be performed using this scale. In other words, the selection of any type of scale of measurement predetermines the future success and minimum performance of that unit, but these limitations also depend on the both the number and type of units included in that level of analysis. This problem has also been described in situations of uncertainty and decision making by James March (March and Olsen, 1976; March 1990), and in information theory by Claude Shannon (1963).

Recent advances in information theory have shown some promise to offering alternatives to these two related problems. Shannon's information entropy is a measure of information content that inherently contains both an upper and lower limit on the complexity of the resulting products. One can use Kolmogorov's theory of complexity to determine the most appropriate level of analysis for any given scale. According to Kolmogorov's theory there exists one explanation for event (possible occurrence or sequence of occurrences) which will explain more of the uncertainty than any other explanation (Cover and Thomas 1992). By using information entropy one can determine what this is, and thus select the scale and scope to match the event (Crutchfield 1992a, 1992b, 1990; Lutz 1990). But it is important to distinguish between simple deterministic systems and complex deterministic systems, because according to Crutchfield (1992a) "there appears to be a way out of the model class discovery dilemma. The answer that hierarchical machine reconstruction gives is to start at the lowest level of representation, the given discrete data, and to build an adaptive series of models within a series of model classes of increasing computational capability until a finite causal model is found. Within each level there is a model-order-estimation inference of optimal models, just as indicated. And there is an induction from a series of approximate models within a lower 'inappropriate' class to the next higher model class" (p. 68)

Shannon's model is based on the amount of information needed to explain or differentiate between the number of choices available. "That information be measured by entropy is, after all, natural when we remember that information, in communication theory, is associated with the amount of freedom of choice we have in constructing messages. Thus for a communication source one can say, ...'This situation is highly organized, it is not characterized by a large degree of randomness or of choice–that is to say, the information (or the entropy) is low" (Shannon and Weaver, 1963, p. 13). If the probability of each outcome is equal then we can say that the selection of any one outcome explains an equal amount of the uncertainty within the decision making process. You could also say that each outcome contributes exactly the same amount of information (or entropy) to the communication.

The amount of entropy contained in a communication becomes important when the total number of outcomes is unknown (uncertainty) or when determining if the information is specific enough (the ambiguity of the message). Thus the information content of an infinite non-repeating string, like that found in the number π (pi), according to Shannon's theory would be as long as the string itself (because no information repeats within the string), and the entropy (or information content) of an infinitely long string of repeating numbers would be trivial. This example also shows that what may be close to impossible to explain at one level (the infinitely long string of the number π), can be easily explained at another level (the written symbol "π"), and that the information entropy explained at these various levels of symbols varies widely. Information entropy, or the information content of any communication that can be transmitted, "is ... approximately the logarithm of the reciprocal probability of a typical long sequence divided by the number of symbols in the sequence" (Shannon and Weaver, 1963, p. 54).

The complexity of decision making behavior can be explained in the same way: "The success at each stage in hierarchical reconstruction is controlled by the amount of given data, since this puts an upper bound on statistical accuracy, and an error threshold, which is largely determined by the observers available computational resources. The goal is to find a finite causal model of minimal size and prediction error while maximizing the extraction of information from the given data" (Crutchfield, 1992a, p. 68).

Kolmogorov's process in creating meaning for complex data streams is to calculate the information entropy (the degree of uncertainty) in every possible chunk from the data stream. In this way Kolmogorov also frees the nonlinear modeling process from a pre-selected data representation, as well as the many problems that *a priori* selection brings. The "chunk" with the lowest entropy will be that data representation that explains the most uncertainty in the given model at that level of analysis. .Shannon's theory allows for the modeling of multiple levels of communication through the idea of calculating the entropy of a string at different levels of analysis, or "morphs" and then calculating the minimum entropy for the string. This allows for a single stream of information to be used in manner that provides adaptive between level representation of the semantic content of a message.

An example of this process that shows how the differences in information entropy can determine meaning in strings based on letter probability and word probability is as follows: if we assume that the total number of possible outcomes (or in this case symbols) are all equally likely, and that these 27 outcomes will be labeled as the letters 'a' through 'z' and space. The probability of a trigram (a string of three symbols) such as "and" given that each outcome is equally likely is 3 times the probability of any one single event (or 3(1/27) = .11). If we look at the probability that the word "and" would occur independent of the letter probabilities then we find that the probability of that word occurring is .028 (27873 occurrences out of a possible 1,000,000 words, Johansson and Hofland, 1989). Thus the lowest entropy we could calculate for the trigram "and" is also the most semantically relevant "morph" or representation.

This example of how information theory can be applied to streams of information to find the level of analysis that provides the most meaning has important implications in how conflict can be resolved within organizations. One way that conflict can exist is between mis-aligned or poorly structured processes or business units. Currently there are two business process restructuring fads that propose alternatives for resolving this type of conflict. The first fad is "business process reengineering", or BPR, (Ham-mer and Champy 1993) which proposes that business should undergo a "fundamental rethinking and radical redesign of business processes to achieve dramatic improvements" (p. 46). It should be obvious that the proposed solution is circular to the degree that the feedback into the system occurs at the same level as the "fundamental rethinking and radical redesign". And traditional BPR tries to push solutions that will result in redesign and feedback occurring at the same levels. The probable result is that the organization of the future will be a horizontally closed system rather than a vertically closed system with control over information access determining what power and authority now determine.

The second fad is the application of open systems theory, or sociotechnical systems (STS), to the design of business processes. Taylor and Felton (1993) recommend that the overall purpose of STS was to "empower" people using a continuum of "information, knowledge, skills and control. These build on one another. Stress is greatest when information is high, but control is low" (p. 7). According to information theory formal control of information will always be low when information is high (and the entropy is evenly distributed). Another consideration of situations when information is high is that the application of knowledge and skills, and ultimately control, is likely to be influenced by highly situational factors—like the size and make up of the group.

The use of metrics of coordination (the modeling of information at several different morphs) can also provide resolutions to several of the problems present in electronic media and group decision support systems. Early theorists in "communication and organization theory" concluded that it was necessary not only to study communication at multiple levels within the organization (Thayer, 1967), but that models of small group behavior must include analysis of communication at the "organizational", "inter-personal" and "intrapersonal" levels. According to Eisenberg (1984) the expectation that organizations of individuals work as "bounded rational" agents is unrealistic because "people in organizations confront multiple situational requirements, develop multiple and often conflicting goals, and respond with communicative strategies which do not always minimize ambiguity, but may nonetheless be effective" (p. 228). One way to study differences between the intra- and interpersonal levels of organizations is to look at organizational behaviors as equivocal communications.

Communication always involves a sender, some content, a receiver and a context. Bavelas, Black, Chovil and Mullett (1990) propose that "the data of communication can be the messages themselves and that the explanation of a message can be sought in the immediate, observable interpersonal situation in which it occurs" (p. 28). John Maynard Smith (1972) has proposed several possible algorithms for modeling communication related to conflict behavior. One of the interesting things about these strategies is that they are not effective until a critical percentage of the population has adopted them—Maynard Smith (1972) found that once the "retaliator" program was predominant within the population none of the other strategies could surpass it.

"One of the principal problems with previous models of the communication process is that they universally neglect the way in which multiple levels of information are exchanged among sender, receiver, and environment" (Targowski and Bowman, 1988, p. 10). In recent years several theories have proposed solutions to the problem of modeling communication between multiple levels. These theories have proposed models for solving the multiple level communication problems in dynamic cognitive environments (Crutchfield, 1992a and b) and for collaborative human-computer communication (Hale, Hurd and Kasper, 1992).

The application of metrics of coordination to group decision support environments can be most helpful in situations of high uncertainty. Within organizational behavior recent studies have shown that the decision making processes (both social and cognitive) in uncertain tasks are recursive (Cohen, March and

Olsen, 1972; Morgan, 1986; Masuch and LaPotin, 1989; March 1990; Beach 1991):. participants return to the same statements and processes until they have been refined or clarified. This recursive process also allows for a model of the outcome likelihood to be developed for group decision making processes. Metrics of coordination can then be used to calculate the entropy of any one outcome within the decision process.

Researchers can also use measures of mutual information and information distance to model such complex behaviors as "brainstorming" and "consensus making." In traditional decision theory the alternatives that will be considered are generated near the start of the processes. These alternatives are then considered and narrowed and some means of consensus making behavior is applied (usually by vote). In uncertain environments researchers found that groups would keep returning to the "brainstorming" stage after they had started evaluation and consensus making (Mintzberg, Rasinghani, and Théorêt 1976; Nutt, 1984)—participants found that their problem space was too confining or that they ran out of alternatives. Using measures of mutual information and information distance (Shannon and Weaver, 1963) the decision making process can be modeled so that participants can monitor the scope of the problem space through the stages of alternative generation and consensus making. In decision making tasks ambiguity is where one level of knowledge or description seems more accurate than a lower level even though that level may have multiple meanings. In information theory ambiguity of this sort is handled by using Kolmogorov's theorems of complexity to model the categorization process. "Kolmogorov [defined] the intrinsic descriptive complexity of an object ... [also known as] the algorithmic (descriptive) complexity of an object to be the length of the shortest binary computer program that describes that object" (Cover and Thomas, 1991, p. 144).

In summary modelers of decision making processes can avoid many of the problems related to designing systems to support decision making behaviors when sufficient complexity exists. Uncertainty can be used to provide designers and computational modelers with problem environments that reflect the cognitive subjectivity and social biases of both human and cognitive agents. Any attempt to explain the variability within any organization must deal with the two problems of scale and scope. And these two problems can be most effectively dealt with in terms of information entropy modeled according to the measures of Kolmogorov's complexity or some other adaptive metric of coordination.

Future Implications for Metrics of Coordination in Organizational Theory:

Organizational theorists have proposed several alternative futures for the interdependencies that people and new technology will share in future work environments. The differences between these theorists can be simplistically stated as those who propose that information will be transformed to adapt to the skills of an individual (Peters, 1992), those who think that the individual will become as specialized as the information they manage (Drucker 1993a and 1993b), and a third who thinks that the workplace of the future will require complex information tools that can track information transfer between variously committed members of teams (Handy 1993 and 1994). While each of the above predictions differ significantly in their implications and assumptions, it is also possible that some combination of all of these alternatives is likely. One of the reasons for this is because all three of the authors deal with the issue of what role information might play in relation to job competence (taking into consideration both the job complexity and the employees skill and experience) without taking into consideration the implications of the scale and scope problems in technological information tools. The issues of scale and scope previously discussed relate directly to how one might try and support these dynamic environments.

Hirschhorn and Mokray (1992) summarize their research on manufacturing automation as it relates to competency and the reciprocity of the role system by offering three possible alternatives for future automated systems: First, that automation can create an environment of reciprocal interdependencies (what they call a "vicious circle") where a breakdown anywhere in the circle of interdependent competencies will result in an imbalance of inputs and outputs as well as a loss of productivity and control. While it is unlikely that once a group descends into such a vicious circle it can get out on its own, there may be the possibility that a change in the metric used to coordinate the groups will result in more cooperation. This possibility has been proposed by Maynard Smith (1972) and summarized by Beniger (1986).

The second alternative that Hirschhorn and Mokray (1992) offer is one where the automation is only partially integrated into the factory. In this scenario a small core of individuals or tasks are automated and the resulting inability to transfer information (in the form of communication, ideas about how to innovate, and scheduling issues) results in "successively smaller and smaller increments in the level of integration and control, so that at some point neither group can stimulate the other group to improve. This is the point where organizational learning stops" (p. 37).

This alternative impacts Handy's (1993, 1994) view of managing performance between team outsiders and insiders based on their commitment to the project. This might only be possible in environments where innovation and learning are not considered important.

And finally, the third option that Hirschhorn and Mokray (1992) propose is that of a "virtuous circle" within which indirect personnel can introduce innovations and direct labor can exert enough control over the system to eliminate bottlenecks before they become problematical. The conditions that make this alternative possible is "a process of sustained reciprocity, whereby information and communication flow freely" (p. 37). Recently research on sustainable cooperation has shown that maintaining cooperative styles in complex environments is more difficult than merely providing sufficient bandwidth to allow for communication. Glance and Huberman (1993, 1994; Huberman and Glance 1993) have shown that an individual agent's tendency to cooperate can depend on the recent history of decisions made in the group, the size of the group, the length of the individuals horizon, and the amount of uncertainty that exists within the system. Further research that will continue exploring this dynamic has been proposed (Fuller 1994). But designers of information tools should not assume that simply increasing the communications bandwidth will be enough to sustain cooperative activity.

References:

Bavelas, J.B., Black, A., Chovil, N. and Mullett, J. (1990). *Equivocal Communication.* London: Sage Publications.

Beach, L.R. (1991). Four revolutions in behavioral decision theory. Conference presentation at the Claremont Applied Leadership Conference. Claremont, CA. Feb, 26, 1991.

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

Callan, J.P., Fawcett, T.E. and Risland, E.L. (1991a). Adaptive case-based reasoning. *Proceedings of the 1991 DARPA Workshop on Case-Based Reasoning.* Washington, D.C.

Callan, J.P., Fawcett, T.E. and Risland, E.L. (1991b). CABOT: An adaptive approach to case-based search. *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence.* Sidney, Australia.

Cohen, M.D.; March, J.G and Olsen, J.P. (1972). A garbage can model of organizational choice. *Administrative Science Quarterly, 17.* P. 1-25.

Cover, T.M.; Gacs, P. and Gray, R.M. (1989). Kolmogorov's contributions to infor-mation theory and algorithmic complex-ity. *The Annals of Probability.* Vol. 17(3). pp. 840-865.

Cover, T.M. and Thomas, J.A. (1991). *Elements of Information Theory.* New York: John Wiley and Sons, Inc.

Crutchfield, J.P. (1992a). Knowledge and meaning: chaos and complexity. in Lam, L. and Naroditsky, V. (Eds.) *Modeling Complex Phenomena.* New York: Springer-Verlag.

Crutchfield, J.P. (1992b). Semantics and thermodynamics. In Casdagli, M. and Eubank, S. (Eds.) *Nonlinear Modeling and Forecasting.* Redwood City, California: Addison-Wesley Publishing Company.

Crutchfield, J.P. (1990). Information and its metric. In L. Lam and H.C. Morris, (Eds.) *Nonlinear Structures in Physical Systems: Pattern Formation, Chaos and Waves.* San Jose State University, San Jose California. November 17-18, 1989.

Drucker, P.F. (1993). *Post-Capitalist Society.* New York: HarperCollins.

Drucker, P.F. (1993). *The Ecological Vision: Reflections on the American Condition.* New Brunswick, New Jersey: Transactions Publsihers.

Eisenberg, E.M. (1984). Ambiguity as strategy in organizational communication. *Communication Monographs,* 51, 227-242.

Fuller, R. (1994). Using information entropy to model a wde variety of complex behaviors within organizations. Published in the proceedings of the 1994 AAAI Spring Symposium on Computational Organ-izational Design. Stanford, California.

Glance, N.S. and Huberman, B.A. (1993). The outbreak of cooperation. *Journal of Mathematical Sociology.* Vol. 17(4). pp. 281-302.

Glance, N.S. and Huberman, B.A. (1994). Organizational fluidity and sustainable cooperation. In K. Carley and M. Prietula (Eds.) *Computational Organizational Theory.* Hilldale, New Jersey: Lawrence Earlbaum Associates.

Greenbaum, J. (1993). A design of one's own: toward participatory design in the United States. In Schuler, D. and Namioka, A. (Eds.) *Participatory Design: Principles and Practices.* Hillsdale, New Jersey: Lawrence Earlbaum and Associates.

Hale, D.P.; Hurd, J.E. & Kasper, G.M. (1991). A knowledge exchange architecture for collaborative human-computer communication. *IEEE Transactions on Systems, Man, and Cybernetics.* 21(3). p. 555-564.

Hammer, M. and Champy, J. (1993). Reengine-ering the Corporation: A Manifesto for Business Revolution. New York: HarperBusiness.

Handy, C. (1994). *The Age of Paradox.* Boston, MA: Harvard Business School Press.

Handy, C. (1993). *Understanding Organizations.* London: Oxford University Press.

Harrison, J.R. and March J.G. (1984). Decision making and postdecision surprises. *Administrative Science Quarterly, 29.* p. 26-42.

Hirschhorn, L. and Mokray, J. (1992). Automation and competency requirements in manufacturing: a case study. In Adler, P.S. (ed.) *Technology and the Future of Work.* London: Oxford University Press.

Holsapple, C.W., Pakath, R., Jakob, V.S. and Zaveri, J.S. (1993). Learning by problem processors: adaptive decision support systems. *Decision Support Systems,* Vol 10, 85-108.

Huberman, B.A. and Glance, N.S. (1993). Diversity and collective action. *Interdisciplinary Approaches to Nonlinear Complex Systems.* Berlin: Springer-Verlag.

Johansson, S. and Hofland, K. (1989). *Frequency analysis of english vocabulary and grammer: based on the LOB corpus.* New York: Oxford Univeristy Press.

Kasperson, C.J. (1985). An explanation of the relationship between performance, decision making, and structure. *Human Relations, 38*(5). 441-456.

Katz, D. and Kahn, R.L. (1978). *The Social Psychology of Organizations,* Second Edition. New York: John Wiley and Sons.

Lutz, J.H. (1990). Category and measure in complexity classes. *1990 SIAM Journal of Computing.* Vol. 19(6). p. 1100-1131.

Malone, T (1988). *Towards a process handbook.* Working paper, Sloan School, MIT.

March, J.G. (1990). How decisions are made. Keynote Address at the 1990 conference on Computer-Supported Cooperative Work. Unpublished document.

March, J.G. and Olsen, J.P. (1976). *Ambiguity and choice in organizations.* Bergen, Norway: Universitetsforlaget.

Masuch, M. and LaPotin, P. (1989). Beyond garbage cans: An AI model of organizational choice. *Administrative Science Quarterly,* 34. p. 38 - 67.

Maynard Smith, J. (1972). *On evolution.* Edinburgh University Press.

Mintzberg, H.; Rasinghani, D. and Théorêt, A. (1976). The structure of "unstructured" decision processes. *Administrative Science Quarterly, 21.* p. 246-275.

Morgan, G. (1986). *Images of Organizations.* Newberry Park, Califor-nia: Sage press.

Norman, D.A. and Draper, S. (1986). *User Centered System Design: New Perspectives on Human Computer Interaction.* Hillsdale, New Jersey: Lawrence Earlbaum and Associates.

Nutt, P.C. (1984). Types of organizational decision processes. *Administrative Science Quarterly, 29.* p. 414-450.

Peters, T. (1992). *Liberation Management: Necessary Disorganization for the Nanosecond Ninties.* New York: Fawcett Columbine.

Shannon, C.E. and Weaver, W. (1963). *The Mathematical Theory of Communication.* Urbana, Illinois: University of Illinois Press.

Spradley, J.P (1979). The Ethnographic Interview. New York: Holt, Rinehart and Winston.

Sycara, K.P. (1993). Machine learning for intelligent support of conflict resolution. *Decision Support Systems,* Vol 10, 121-136.

Targowski, A.S. and Bowman, J.P. (1988). The layer-based, pragmatic model of the communication process. *The Journal of Business Communication.* Vol. 25. p. 5-24.

Taylor, J.C. and Felton, D.F. (1993). *Performance by Design: Sociotechnical Systems in North America.* Englewood Cliffs, New Jersey: Prentice Hall.

Thayer, L. (1967). Communication and organizational theory. In F. E. X. Dance (Ed.) *Human Communication Theory: Orginal Essays.* New York: Holt, Rinehart and Winston, Inc.

Weick, K.E. (1979). *The social psychology of organizations.* Second edition. New York: Random House.