

Bridging the Gap: On the Collaboration between Symbolic Interactionism and Distributed Artificial Intelligence in the Field of Multi-Agent Systems Research

Jörg Strübing
Free University of Berlin

This article discusses the relationship in the U.S. between current symbolic interactionism and computer sciences—specifically, distributed artificial intelligence (DAI). The general thesis is twofold. First, current interactionist approaches to organization, science, and technology show a special affinity to goals and problems of DAI research, and in research style, methods, and theoretical concepts, symbolic interactionism can provide useful suggestions in the design of DAI systems. Second, a good way to analyze the relationship between computer sciences and symbolic interactionism is reflexive of theoretical concepts provided by interactionist approaches. In this sense, DAI is a “going concern” which extends across various fields and intersecting social worlds connected through a set of conceptual “boundary objects.” It is concluded that the interaction between technology and sociological thought must go beyond a mere exchange of ideas. What is required is continual, hands-on, trans-disciplinary collaboration.

INTRODUCTION

The relatively young research domain of distributed artificial intelligence (DAI) and the somewhat older discipline of sociology have been sending out feelers to each other for some time. Distributed artificial intelligence means modeling and designing distributed computer systems to serve the needs of large technical infrastructures with real-world applications. This can be seen as a kind of maverick branch of artificial intelligence (AI), since it seeks to abandon efforts to reproduce or emulate individual intelligence. Instead, DAI—or more precisely its special branch of multi-agent systems (MAS) research—seeks

Direct all correspondence to Jörg Strübing, Institute for Sociology, Free University of Berlin, Babelsberger Str. 14–16, D-10715 Berlin, Germany; e-mail: jstrueb@zedat.fu-berlin.de.

to combine people and computers in a collective, real-world approach.¹ Thus, distributed artificial intelligence looks to sociology for ideas that could serve as a basis for better models of collective problem-solving. Sociology, the more cautious partner (Rammert 1995, p. 17), has shown some interest in DAI's proposed efforts to simulate social phenomena. I will discuss the various ways in which sociology may contribute to MAS research as well as possible forms of cooperation and exchange between the fields. But before doing so, it would be useful to have a brief look at (1) the questions DAI puts to sociology or to an everyday understanding of social phenomena, (2) the goals DAI is pursuing, and (3) the problems which it is unable to solve within the framework of technical sciences alone.

In discussing the potential benefits from DAI's interest in sociology, the question arises as to which of the various "schools" of sociology has the most to offer. Sociology, like other scientific disciplines, is far from homogeneous. Current symbolic interactionism (SI) is the branch of sociology which is especially close to the goals and problems of MAS research. Moreover, in the areas of research style, methods, and theoretical concepts, it has especially useful suggestions to make toward the designing of MAS. Current SI can provide empirically-grounded theoretical concepts that describe social phenomena, and these concepts are transferable to MAS. Its systematic approach to empirical research methods also offers a model for both the process of designing MAS and for modeling knowledge-generation processes. In this article, I will present some interactionist concepts and discuss their value for MAS research. DAI and SI have been involved with each other for a long time in the U.S. It is not only that interactionist concepts have deeply influenced DAI models; a good number of these same concepts have been developed and refined in empirical research on different aspects of distributed artificial intelligence in socio-technical environments.

Thomas Malsch (1997) asks the following questions: Which paths of migration are being used in transferring sociological knowledge to DAI? What are their constitutive properties? My position is that the paths of continuous cooperation and doing research together on projects of mutual interest is especially fruitful. It is not sufficient simply to transfer social metaphors to DAI. The more promising approach is to establish ongoing cooperative structures ("going concerns" as Hughes [1971] would call them) aligned by and centered around shared research questions and tasks, acting as so-called "boundary objects" (Star 1989). Establishing and maintaining such "discipline-linking" cooperation is essential if DAI and sociology are to come closer together.

In this article, the term "current SI" refers to one recent direction in interactionist research and theory, from which numerous empirical studies of organization, science, and technology have emerged. (For other current approaches in SI, see Denzin 1992; Ellis and Flaherty 1992; and Prus 1996). These studies have built on Strauss's and Becker's extension of Blumer's work on processes of identity formation and collective behavior as well as on the theoretical framework elaborated by Park and Hughes in their empirical work on communities, professionalization, and the division of labor (Fisher and Strauss 1978, 1979; Star 1996, p. 298; Strübing 1997).

A short sketch of some central questions in MAS research is followed by a description of the most important conceptual links between SI and MAS. The second section discusses

Hughes's "ecological approach," while the third section goes into the possible usefulness of the *negotiated order approach* and *social worlds theory* in the modeling of cooperation and mediation in MAS. The fourth section links this whole discussion to the "scientific community metaphor," a common focus in MAS, while the fifth section deals with the question of the multi-agent systems' ability to interact. In order to support my thesis in regard to the importance of continual and practical cooperation between sociology and DAI, the sixth section analyzes the cooperation between current SI and American DAI research. The concluding remarks present some arguments on the relationship between current SI and the micro-macro question in sociology.

GOALS AND PROBLEMS OF RESEARCH IN DISTRIBUTED ARTIFICIAL INTELLIGENCE

Research in DAI began in response to the difficulties experienced with traditional AI models in designing problem-solving systems useful in practice (Hayes-Roth 1980). DAI's central thesis is as follows: It is not feasible to model real-world problem-solving processes (i.e., the emergence of intelligence [Star 1989, p. 41]) as long as AI research clings to the idea of a single "autonomous" and omnipotent machine which develops solutions simply by making rational use of all available resources. This is because, in the real social world, cognition and intelligence (i.e., the ability to solve problems) are distributed phenomena which can only be addressed by the joint efforts of various actors (be they machines, humans, or other living things). When trying to model these indeterminate processes in technical systems, concepts like omnipotence, autonomy, or universal rationality are inapplicable. Instead, we should conceive of models in which technical agents, with a limited degree of autonomy and some ability to solve problems, work collectively in solving a problem.²

AI set out to model, predict, and even simulate the actions and thought processes of people within a social context, but it kept as much of this social context as possible outside of the models. Systems of knowledge are becoming increasingly embedded within larger networks of humans and machines, where problems are distributed and open-ended (Gasser 1991, p. 108; Hewitt 1985, 1986).

Sociologists know that this model-building is not a simple task. The models developed in DAI have been quite varied, as have the modeling approaches behind them (Schulz-Schaeffer 1993; Müller 1993; Gasser 1992). However, a number of common problems and concerns can be distinguished. Three "basic problems" of DAI are listed by Müller (1993, p. 11): (1) the "agent-construction problem" (what properties should the entities of a distributed system have?); (2) the "society-construction problem" (how should these entities or agents interact?); and (3) the "social-application problem" (evidence of real-world problem-solving capacity of DAI systems and profiles for useful applications). Looking at it in terms of social reality, what we have here are the central areas of investigation in sociology:

- What are actors?

- How do organizations and societies work?
- What do different types of organizations and societies accomplish?

For DAI, defining what agents are or should be means distinguishing between (1) cases where systems are already in place to which certain features or properties would need to be added in order for it to be integrated into problem-solving networks (i.e., adding communication protocols) and (2) cases where all agents need to be completely redefined or invented. In the latter, the question arises as to what extent single agents need support from other agents in order to solve problems. Likewise, it must be determined whether agents are capable of solving problems and taking action by themselves or whether they can only do so in cooperation with other agents. (For a typology of agents, see Nwana 1996.) These questions can only be answered when one takes into consideration the particular organizational or social model involved. The answers depend on both the coordination and cooperation concepts at the agents' disposal and on the limits of the system as a whole (rules of participation, openness or closedness of nets, etc.).

One can also approach the difficulties in modeling DAI systems from the perspective of the problem to be solved, as does Gasser (1992, p. 11): First, the problem must be translated into a suitable interactive language or into a mechanism of representation suitable for DAI. Second, rules and procedures are designed for decomposing and allocating the problem. Simultaneously, the way back from the problem-solving activities of single agents to comprehensive solutions which are presentable to the "world outside" needs to be modeled. In finding solutions to these questions, one is faced with at least two important obstacles.

One is *the heterogeneity of agents*. If getting the most use out of DAI systems means combining the problem-solving capacity of different agents in order to achieve synergetic effects, then the participating agents must differ in the ways they approach the problem. However, that would mean that there is no such thing as a commonly shared "global perspective" (Gasser 1992, p. 13). Therefore, coordinating the distributed system through a central authority does not sound very promising. The other is *the openness of systems*. Since one of the most striking advantages of DAI systems is their ability to handle contingent events in a flexible manner, it follows that agents cannot have at their disposal a complete set of rules and knowledge *ex ante*. Thus, it makes sense to design such distributed systems as open networks in which learning is possible and agents' roles are not pre-assigned. Instead, the system should allow for flexibility in deciding whether (1) further agents should be integrated, (2) if new agents are to be generated, or (3) an agent is to be dismissed. All of this, however, makes coping with contingency much more difficult, as one can be sure of neither the number and the type of agents involved, nor of their mutual perspectives on the problem.

The result of all this is that the desired emergence of solutions to problems is accompanied by highly unpredictable system behavior. Hewitt coined the term "unbounded non-determinism" in order to label this basic property of open systems (Hewitt 1985; see also Star 1996, p. 304). At any given moment in the process, there is no *a priori* certainty as to which agents are involved, what their aims are, what the consequences are for the system as a

whole or, most importantly, which synergetic effects might result from the interaction of agents.

ECOLOGICAL APPROACH AND MULTI-AGENT SYSTEMS

In view of the above goals and problems in MAS research, two questions arise: (1) What are SI's strengths as a sociological source of inspiration, and (2) how does this process of inspiration proceed? First, we should bear in mind that DAI and SI have been working together for some time. For at least 15 years, important researchers from the American DAI community—such as Les Gasser, Carl Hewitt, and Alan Bond—have built extensively on the approaches to organization, science, and technology developed by Anselm Strauss and his colleagues at the University of California, San Francisco, over several decades, and by Elihu Gerson and colleagues from the Tremont Research Institute in San Francisco in the early eighties. In 1988, when DAI was beginning to establish itself as a field of research, a circle of interactionist and pragmatist sociologists and computer scientists produced a position paper calling for increased institutionalized cooperation between the computer sciences and the humanities/social sciences (Bendifallah et al. 1988). In keeping with this proposal, Susan Leigh Star, a former research associate at the Tremont Institute and a symbolic interactionist student of Strauss, spent several years teaching and doing research in a computer science department, and now is in an information science department.

One of the most challenging problems in constructing multi-agent-systems is to find suitable concepts that can deal with both the internal constraints of software and the various qualities of social processes. Carl Hewitt (1985, 1986) stated that, in order to solve real-world problems, MAS needs to adequately represent the structures of the real world. Real-life problem contexts are characterized by “openness”—that is, the absence of both spatial or temporal confinement and of a central decision-maker. Bond and Gasser (1988, p. 8) summarize Hewitt's statement as follows: “Hewitt notes that many real-world distributed systems (e.g., office work . . . large distributed data bases . . .) exhibit characteristics of (1) mutual inconsistency in knowledge or belief; (2) bounded influences or ‘arms-length relationships’ among components; (3) asynchronous operation; (4) concurrency; and (5) decentralized control.”

Hewitt's “open systems approach” is very similar to the “ecological approach” developed several decades earlier by E. C. Hughes (1971).³ Hughes's general idea is that organizations and other social institutions must be thought of as “*going concerns*” which are constituted by “groups of people sufficiently committed to something to act in concert over time” (Clarke and Gerson 1992, p. 187).⁴ A going concern represents a set of (often heterogeneous) ecologies, which vary with the perspectives of the various actors involved.

The basic idea behind this approach is that social aggregations (groups, organizations, or societies as a whole) are based on temporary and permanent interactional contexts. Instead of being the “glue” that keeps social institutions together, formal frames are secondary phenomena maintained and modified by social interaction. Mead (1934, p. 327f) already followed this track when he stated, “That is what makes communication in the significant

sense the organizing process in the community. It is not simply a process of transferring abstract symbols; it is always a gesture in a social act which calls out in the individual himself the tendency to do the same act that is called out in others." However, while Mead developed this perspective within the context of social philosophy, it was Hughes who demonstrated empirically that Mead's "universes of discourse" *are* a decisive driving force in organizations.

In doing so, Hughes shifted the argument slightly in another direction. While Mead's use of the term "universes of discourse" fostered the misunderstanding that communication instead of social action was the cohesive force, Hughes's focus was explicitly social interaction and work.⁵ At about this same time, Shibutani (1955) was working on the concept of "reference groups" as perspectives. This approach also denied the dominant role of formal membership criteria. In contrast to Hughes, Shibutani (1955, p. 566) stressed "effective communication" instead of mutually related action.

Looking at sociology through the eyes of DAI researchers, a communication-based concept might seem more promising than a focus on action. Strictly speaking, relations between different components of a distributed system must be established via (technical) communication. However, it is not the fact of communication which constitutes a social relationship, but rather content and the social act of communicating itself. The reference to actions occurring elsewhere constitutes the social context, or, in Hughes's words, a "going concern." These actions can (but need not) refer intentionally to each other. The same applies to DAI systems. We can think of the technical components of a distributed system as performing actions which might be induced through an interactively constituted going concern (e.g., solving a defined problem), but each action might just as well be induced through other contexts. In every case, the relevance of a certain activity to a particular going concern must be determined in discourse.

When we consider (1) that technical agents are able to act simultaneously in different problem-solving processes and (2) that the effectiveness of multi-agent systems derives from the integration of different types of agents, we come up against a problem very familiar to sociologists: The differences between agents correspond with a *multiplicity of perspectives* effective in the process. As with social interaction, in MAS, we cannot expect all acting entities to share a common perspective on the process or its goal. However, this multiplicity does not necessarily result in conflict. Here, multiplicity means that perspective is determined by situation and position. Where you sit is where you stand, as the saying goes. Divergent viewpoints may result in different interpretations of a problem, in heterogeneous problem-solving strategies, and in divergent methodology. This might sound problematic with respect to coordination and cooperation, but it is precisely with the help of this divergence that DAI researchers expect creative solutions.

However, a problem arises when we try to prove an analytically accurate representation of the structure of variably integrated heterogeneous entities and their interplay. Hughes's ecology of institutions tries to solve the problem through the methodological guideline of a triangulation of perspectives: A sociologically satisfactory analysis of a going concern requires (1) that data be collected on the relevant ecologies of the participating actors, including their alteration over time, and (2) that the manner in which these ecologies relate

to each other be made explicit (Hughes 1971, p. 19; Star and Griesemer 1989, p. 389). For the social process itself this means, that success depends on the solutions that are found in relating the divergent perspectives of the various actors to each other to form a more or less coherent problem-solving process. With respect to this process of integration, SI is very skeptical about the idea of a central directing agent, at least insofar as it concerns average social processes. The more critical DAI researchers share this skepticism (Bond and Gasser 1988, p. 8; Galliers 1990, p. 36).

MEDIATION, CONTROL, AND COORDINATION IN OPEN SYSTEMS

Clearly, a hierarchical, mono-focal perspective representing only a particular actor's viewpoint will not lead to a sound representation of the social process in question. Around the next corner, we face the question of how and through whom these heterogeneous perspectives become practically integrated. No going concern survives without alignment (Strauss 1985).

Hewitt's actor model (1977; see also Agha 1986), widespread among MAS researchers, regards every component of a system as a closed unit (i.e., every agent is a black box to other agents). Exchange between agents takes place via communication channels. In contrast to lines of command and control, every individual agent has to decide which messages to respond to and act on. Actions can include sending a new message to certain other agents, creating new agents, and integrating them into the problem-solving context. Here, hierarchical, command-based structures are not efficient in coordinating the problem-solving context. There is no superior authority who (1) knows enough about what is going on in every single agent and (2) is able to determine which agent it would make sense to integrate during a certain phase of a problem-solving process.

In seeking adequate concepts for the coordination and control that takes place in open systems, MAS research could benefit from cooperating with sociology—as long as these concepts are computable (see Gasser 1991).⁶ The ecological approach offers little more than an indication of the role of commitment for an understanding of cooperation in open ecologies. However, the empirical studies of Strauss and his associates have gone far beyond Hughes's approach. One aspect of their theoretical framework is the "negotiated order approach" (Strauss et al. 1963; Strauss 1978a; Maines and Charlton 1985). Although Strauss does not claim that organizations or other social constellations emerge solely from processes of negotiation, he *does* claim that every social order and every organization is necessarily based on negotiation and, moreover, that negotiations are essential not only to the emergence of organizations but to their maintenance as well (Strauss 1978a, p. 235). A remarkable number of MAS researchers have made use of this concept in order to develop solutions to the coordination problem (Bond 1991; Davis and Smith 1983; Durfee, Lesser and Corkill 1988; Adler et al. 1989; Hewitt 1986; Gasser et al. 1989).

In real world negotiations, we do not have completely autonomous subjects acting in a structureless vacuum. Instead, both the negotiation process itself as well as the negotiating parties are part of other social and material contexts. This is the relevance of Becker's

(1960) notion of "side bets." Social worlds theory also describes processes of negotiation as taking place between actors representing different social worlds: "groups with shared commitments to certain activities, sharing resources of many kinds to achieve their goals, and building shared ideologies about how to get about their business" (Clarke 1991, p. 131). While each social world produces unique structures of relevance, it is evident that negotiating actors represent different perspectives and act in accordance with other frames of reference which are not entirely present in the negotiation (Strauss 1978b; Clarke 1991). Following Strauss and his associates, the field of action, where negotiations between social worlds (or between social subworlds; see Strauss 1984) take place, might have the structure of an "arena" where the field covers a widespread set of controversies and representatives of different social worlds are involved (Clarke 1991, p. 133). However, there are also multi-world negotiations below this meso/macro level of "arenas." These constitute a going concern.

Thinking of MAS in terms of social worlds and negotiated orders allows us to perceive the continuous changes which networks undergo. New agents get involved; others leave the "on-going" concern; the amount of knowledge they have is subject to constant change. Cooperation among agents from similar structures and contexts (i.e., those who belong to the same "technical world") can be expected to proceed more smoothly. On the other hand, their range of perspectives is limited and so is their capacity to produce emergent solutions. Heterogeneous teams of agents are exactly the opposite: they are complementary in their problem-solving ability, but find it more difficult to achieve commonality for collective operations (the problems of translation and alignment).

Against this background, it becomes clear that Gasser's definition of organizations as "a particular set of settled and unsettled questions about belief and action through which agents view other agents" (Gasser et al. 1989, p. 58; see also Bendifallah et al. 1988) is anything but an expression of subjectivism. Gasser and his colleagues explicitly refer to the negotiated order approach and the social worlds theory (i.e., to the social construction of actors' capacity to act). Those who advocate current SI as well as DAI researchers like Gasser and Hewitt draw inspiration from the insight that social structures can only be experienced in and changed through social interaction.

With this interpretation of organized problem-solving as negotiation among representatives of social worlds in arenas of mutual interest, the process itself still remains insufficiently determined. How is effective problem-solving achieved? Under what conditions do actors in heterogeneous settings cooperate? Both Hewitt (1986) as MAS researcher and Gerson and Star (1986) as sociologists call this the problem of "due process." Thus, Star (1989, p. 42) asks, "[I]n combining or collecting evidence from different viewpoints (or heterogeneous nodes), how do you decide that sufficient, reliable and fair amounts of evidence have been collected? Who or what does the reconciling, according to what set of rules?" A kind of test is needed to determine the appropriateness of rules and procedures in open decision processes.

These questions are not only central to MAS research, but are also high on the list of sociologists doing research on science and technology (Heintz 1993b; Callon 1994; Clarke and Gerson 1992). Many of these scholars have built on the preliminary work of Strauss

and his associates (1985; Strauss 1985, 1988; Gerson 1983), who showed empirically how the organizing of work processes is itself an integral part of the work. Strauss conceived of work as a combination of individual tasks that form “arcs of work.” How the tasks are combined and the way arcs of work are coordinated with “lines of work” (i.e., the performance of a specialized skill or discipline) are questions to be negotiated. In this fashion, Strauss identifies the problem of contingency inherent in work processes. He argues that an “arc of work” can be planned bit by bit, but it can never be foreseen or planned as a whole. Inevitably contingent events occur which require an interactive realignment of task performance and task sequences. Thus, only in retrospect do we get a complete picture of a work process (Strauss 1985, p. 4). When applied to DAI-models, this idea has significant implications: Trying to model a completely formalized and pre-determined problem-solving process does not seem very promising. Instead, we need to implement something like “process intelligence.”

Recent studies have developed Strauss’s approach much further and made it applicable to science and technology. As a result of her empirical studies of the “molecular biological bandwagon” in cancer research, Fujimura (1987) suggests that we focus on the transformation of a hitherto unsolved problem into a “doable problem.” Doable problems are processes of interaction in which standard operation procedures are established which break complex (and seemingly unique) tasks into aligned pieces of routine work. This interactional process of establishing these procedures is what Strauss termed “articulation work” (Strauss 1985, 1988). The important point is that in producing “doability” the respective interpretations of the problem at hand are made compatible but not necessarily identical. Recent interactionist studies of science and technology have revealed that these interactive processes of alignment and structuring are not exclusively determined in communicative acts, but are also developed in material practice (see Clarke and Fujimura 1992). It is not the inspired idea alone that makes problems “doable,” but rather a chain of mutually integrated and practical actions.

This same viewpoint is expressed in Star and Griesemer’s (1989; Star 1989) concept of “boundary objects.” They stress that under conditions of heterogeneous and spatially distributed cooperation, an important part of the articulation process and interactional alignment is performed via mediating objects. Explicit consent for procedures and goals is not required in order to establish a practical going concern. Instead, actors in a field of action can align their mutual activities by creating boundary objects (artifacts, ideas, experiences, animate or inanimate objects found in nature, etc.). These efforts are successful as long as these objects prove both stable enough to provide a common identity across sites and flexible enough to allow for all parties to apply their particular perspectives to these objects—without the object breaking up into a number of incoherent parts (Star and Griesemer 1989, p. 393).

Other interactionists have made similar proposals for processes of mediation and alignment in heterogeneous settings (Fujimura 1988; Henderson 1995). Of central importance in all of these efforts is the attempt to explore how projects in science and technology succeed despite increasing heterogeneity in actor constellations and the growing complexity in cooperative structures. Star writes that the scientists she observed not only lacked informa-

tion on the work style of the other actors, they did not even employ the same methods, analytical units, or data models. Moreover, their cooperation was based on divergent goals, time frames, and reference groups. Nonetheless, the result was often a successfully distributed process of problem-solving and not chaos. Consensus was not required (Star 1989, p. 46).

THE SCIENTIFIC COMMUNITY METAPHOR

While sociology deals with the problem of how to discover and analyze these (“due”) processes, MAS research aims to model similar processes. An early MAS model for this type of cooperation between distributed technical agents was the “scientific community metaphor” introduced by Hewitt and Kornfeld (1988; Hewitt 1977). This model, inspired by both the Kuhnian philosophy of science and some interactionist ideas, suggests that we view distributed systems as social communities of experts dealing with a research question of common interest. However, some authors (e.g., Schulz-Schaeffer 1993) conclude that through the use of this metaphor, the “benevolent assumption” (as cited in Galliers 1990; v on Martial 1990) found its way into MAS discussions. This is the assumption that all participating actors share “goodwill” in approaching a common goal. As explained above, the “benevolent assumption” is not state-of-the-art in MAS research.⁷ Nevertheless, there are good reasons to think about the scientific community metaphor and its relation to MAS. Obviously, the idea that scientific communities might have something in common with shared goodwill—as posed by some critics—is based more on an obsolete common-sense image of scientific communities than on its meaning in terms of the philosophy of science during the last three decades. Currently, another feature of scientific communities has come to the fore. Gasser describes it as the interplay of local construction (of theories, solutions, or techniques) and the communicating of these constructions among problem solvers and scientists located in other local contexts who test, reshape, or generally question proposed models and designs (Gasser 1992, p. 14). Gasser leaves homogeneity and like-mindedness completely out of the picture. In relation to technical systems, this means thinking of nodes in networks as being able to propose (parts of) solutions, which do not automatically become valid, but which are subject to further tests, changes, and supplements by other competent nodes. Furthermore, the problem-solving process should be documented and published within the net to allow for a re-accessing of archived solutions. The whole module is structured as a dynamic network, similar to, for instance, e-mail discussion groups on the internet. Thus, choosing the scientific community metaphor as a basic analogy by no means denies that there is asynchronous communication, an imbalance of power, practices of non-disclosure, or large differences in the availability of resources. Speaking in terms of SI, the homogeneous group of like-minded experts—as addressed in what I termed an obsolete common-sense understanding of scientific communities—has turned into an arena of heterogeneous actors belonging to different social worlds. It is precisely from this perspective that it makes sense to think about the analogy of scientific communities and MAS.⁸

MACHINES AND INTERACTION—A CONTRADICTION?

Substantiating the affinity of SI and MAS must go beyond simply listing analogies between concepts used in the two fields. We need to discuss the possible or actual application of basic theoretical SI notions in the domain of MAS. As the name suggests, SI is based on a specific view of interaction which determines who the actors are: individual actors and, even more so, institutions which have always been socially constructed. The individuals exist as intelligent problem-solving beings through their participation in processes of communication and interaction. All action takes place within the context of interaction. Mead's (1934) concept of symbolic interaction referred to human behavior, and it should be noted that the idea of "significance" was meant to distinguish humans from animals. It was not meant as a criterion for distinguishing between human and technical agents or in building sociologically inspired models of technical communication.⁹ In attempting the latter, we have to face the following questions: Are technical agents able to produce significant gestures which do not lead those they address to a determined response but rather to an interpretation of its meaning? And could such a gesture-producing agent recognize in the resulting behavior of other agents a *result* of its gesture (i.e., would they be able to experience)? Finally, will it be advantageous for MAS research to operate within symbolic interactionist models of action?

To begin with, it should be noted that MAS research already *does* operate with these models. For instance, among the essential properties of actor models, Bond and Gasser (1988, p. 8) list "control structures as patterns rather than as sequential choices among alternatives." The model of an exchange of messages substitutes for signal-action-signal structures most commonly used in technical artifacts. Reflexivity is the new element. In the following stimulus-response-based communication, "A to B: case X happened, start action X_1 / B to A: X_1 started, result: status Y," it is determined *a priori* what B is to do if case X happens. However, in a reflexive message-exchange model, we would design the communication process in such a way that incoming messages are interpreted by actors in light of their respective goals or structures of relevance. Furthermore, decisions on how to proceed would be made by the receiving actor. This makes sense only if the sending agent recognizes the other's action as a response to the initial message. That is, the initial message *acquires* significance. The sum of these registered responses would give each actor an internal representation of the reactions to expect in response to significant messages. Thus, we can think of this as an attitude of expectation, which at the same time implies the possibility of disappointment and/or of novelty.

In MAS research, the concept of agents as partly autonomous and partly able to solve problems could be represented in an analogous way. Every single agent would need the interpretative reactions of the other agents to understand the meaning of their own actions. The recent claim by some interactionists that technical artifacts are capable of action, far from being eccentric, turns out to be both obvious and productive (Clarke and Gerson 1992).¹⁰ One might still object that processes of human identity formation on the one hand and the manipulation of bits and bytes on the other cannot play in the same league (i.e., that they differ ontologically). Looking at things pragmatically, however, the question is mis-

guided.¹¹ Since the pragmatist's view is less concerned with ontological metaphysics than with practical consequences of action, interactionists are interested in the question of whether regarding machine systems as networks of acting agents makes a significant difference and, if so, what the difference might be. Moreover, in recent SI, the more far-reaching thesis denies the appropriateness of analyzing technical artifacts—be they single machines or DAI systems—separate from their social context of development and usage.¹² What a technical artifact is, what it achieves, and what practical consequences it produces can only be decided in light of its social context (Star 1989).

Behind all this we find Thomas and Thomas's famous definition of the situation, which was possibly the first instance where sociology benefited from pragmatism. They state, "If men define situations as real, they are real in their consequences" (Thomas and Thomas 1928, p. 572).¹³ In any given situation, there is no *a priori* difference between human and non-human actors, as long as they themselves do not experience and define such a difference. This refers not simply to the output of a situation—whether, for example, a machine reacts the same way as a human being in a certain situation; what is also involved is Mead's idea of self-significance. In specifying self-significance as constitutive of the identity of actors, Mead was not speaking in ontological terms, but was referring to layers of experience without which further acting is impossible.

A second approach to the question of machines' capacity to act can be observed in recent SI studies. In general sociology as well as SI, it is widely accepted and does not require much effort to construct collective actors and to attribute to them the ability to act (Strauss 1993, p. 23). However, SI often aims to decompose (not to reduce!) collective actors into their components and to deconstruct organizational processes into their single interactions. (This is, for instance, the meaning of the negotiated order approach.) Thus, from the perspective of SI, it is not very difficult to accept the idea of non-human actors: In fact, behind every MAS, as with any light bulb, we find human labor. But, as Star (1991) and Star and Strauss (1997) put it, this is "invisible work" which is part of the artifact as such and comes into play again when it is used. At the same time, for users working with these artifacts, there is no access to the full social and political dimensions of the context within which the artifacts have been designed and produced. They are black boxes.

INTERACTIONAL ALIGNMENT VS. MIGRATION OF METAPHORS: ON THE MODE OF COOPERATION BETWEEN DAI AND CURRENT SI IN THE U.S.

Thomas Malsch (1997, p. 16) proposes that we regard the interaction going on between sociology and DAI as a "migration of metaphors." He means that the type of transfer taking place cannot be described as mere translation. Rather, "technological innovation can come about when concepts are removed from their original frame of reference, i.e., human society, and are transferred as social metaphors, via several intermediary steps, first to computer science concepts and formulae and then, finally, to technical artifacts" (Malsch 1997, p. 16). A direct analogy between society and artifact would not be tenable because, in DAI

research itself, there is an inherent “tension between social reference and computer reference.”

I am especially interested in these intermediary steps. My *thesis* is that an important part of the American MAS community, which Malsch rates as world leading (1997, p. 15), can be characterized as a well organized going concern, which is successfully mediating between SI and MAS by establishing a number of boundary objects. It is not simply that sociological concepts have been transferred to computer sciences via these boundary objects. Shared ideas, models, and practices relevant in both domains have also developed jointly. The crucial point is that, in order to correctly understand the route from a social metaphor to a technical artifact modeled on this metaphor, we must include these organizational structures in the analysis, or, more precisely, those processes of distributed organizing and problem-solving which made it possible to link metaphor and artifact.

It is important to note that the various authors of the “white paper” on DAI research (Bendifallah et al. 1988) regularly cooperate with each other in several other contexts beyond the field of DAI. That is, interactionist sociologists in the field of science and technology studies (STS) do not provide input into DAI research on only a sporadic basis as, for instance, contributors to DAI conferences or DAI books. Instead, there is a continual process of interaction in joint research, publications and conferences, the exchange of manuscripts, and so forth.¹⁴ What is the reason for this mutual involvement? How have these connections across the borders of not only disciplines but cultures (as Snow 1993 put it) been established? How has the arena of interactionist STS research and DAI been maintained and stabilized? In my view, there are two factors relevant in answering these questions: first, a shared interest in a set of research questions which are equally relevant to both domains (though with different connotations), and second, a common philosophical ground in pragmatism and a concomitant approach in research.

Shared Research Questions

What drives computer scientists in the domain of MAS, such as Hewitt or Gasser, and interactionist sociologists, like Gerson and Star, are what they call “due process” (i.e., processes of decision-making and problem-solving in heterogeneous settings, where there is no predetermining logic of decision and no universal criteria of relevance). This area of inquiry is as relevant to research in the history of science (Star and Griesemer 1989) as it is to the development of MAS models. Star (1989, p. 37) expresses this link in her definition: “Due process means incorporating different viewpoints for decision-making in a fair and flexible manner. It is the analog of the frame problem in artificial intelligence.” However, the ways MAS and sociology look at this problem differ considerably. While sociology came across due processes in attempting to examine and explain modern societies, MAS is interested in due process as part of its effort to construct suitable models.

Similarly, the problem of open networks can be understood as a common focus of both domains. Coinciding with MAS’s emancipation from traditional AI’s methodological individualism, and with the development of multi-agent-networks, recent SI research has been focusing less on situated interaction processes. Instead, it has integrated overarching struc-

tural conditions with the practices of local actors in an effort to establish a network view of these processes (see Bowker and Star 1996; Star and Ruhleder 1996).

Another shared concern between MAS research and SI is the rejection of determinism and functionalism in scientific analysis. Thus, Gasser (1991) criticizes deductive logic as an inadequate basis for scientific reasoning; Hewitt (1985) denounces the closed-world of traditional AI as insufficient for real-world applications; and Star (1996, p. 303) challenges functionalism by stressing a "concept of undetermined conditions." They all point in the same direction, and, moreover, they gain some mutual reinforcement. Here (e.g., in Star 1996, p. 304), we can see that arguments and concepts are heading in both directions and that MAS findings support interactionist research as much as the reverse.

A good example of the primacy of practical cooperation over a mere exchange of ideas can be seen in Star's proposal for a suitable evaluation procedure for MAS. She claims that MAS programs have always been socio-technical. They emerge from a process of social practice, and they develop in a social context of use. Development and use are inseparable because what is reified in the development of systems becomes social once again when they are used. In the process, the original intent of the artifact is modified.¹⁵ Consequently, assessing the quality of a socio-technical system cannot be limited to the technical artifact itself, nor can these evaluations be conducted *ex post*. As Star stated, as long as we refer only to the artifact, we end up with the same shortcomings we found in the Turing test.¹⁶ This test, which was developed by computer scientist Alan Turing more as a theoretical model than as a practical evaluation procedure, was criticized by Star for postulating a closed test universe, in which the test is performed by individuals instead of communities, and where computers are regarded as universal machines limited only by storage capacity and processor speed (Star 1989, p. 39). Star quotes Hewitt when pointing out that no enlargement of technical capacity can overcome the problem of operation in open systems.

Thus, in order to adequately validate the quality of MAS (i.e., to determine whether problem-solving capacity has emerged) she proposes a "Durkheim test." This test would take into account the fact that in distributed open systems it is feasible neither to test the whole system nor to find out about the intelligence of a system by testing only one of its components. Star suggests that in dealing with the problem of evaluation under conditions of emerging contingency, we should evaluate the complete process of a participatory system's development. In this way, the question of emerging intelligence would become a question of social visibility (Star 1989, p. 41). That is, both the interactive process of designing the system and the social meaning the system acquires through being used would be included. Her proposal basically draws on Hughes's methodological statement that we cannot adequately analyze an institution or an organizational setting without access to data revealing the inner processes and the multitude of perspectives relevant in that going concern (Hughes 1971, p. 19).

Common Research Style

Like other interactionist researchers in the fields of technology, science, work, and organization, Star has built on the research style of "grounded theory."¹⁷ Denying the validity

of much traditional social research, grounded theory specifies practical usability of results (especially for the field in question) as the central measuring rod for good research (Corbin and Strauss 1990, p. 422). Thus, grounded theory points the way to ethnographic field work designs, which both triangulate data from different perspectives and re-address the results to the people acting in the field. This would imply a type of action research, which aims at both learning about as well as changing social reality (exactly what Star proposes with the Durkheim test).

It is interesting to see the parallel in one basic position shared by grounded theory and MAS: "There is no central arbiter of truth," as Kornfeld and Hewitt noted (1988, p. 312)—simultaneously addressing MAS-models and scientific communities. This same initial consideration, which leads to decentralized decision-making models in MAS research, causes grounded theory to question the appropriateness of traditional quality standards in empirical research. The absence of a universal rationality or a general criterion for truth requires us to re-contextualize every decision-making process.

We can see other parallels if we direct our attention to the value both MAS research and adherents of grounded theory attach to common-sense knowledge. For instance, Kornfeld and Hewitt (1988, p. 314) make the following observation: "An important source of proposals are very general methods coming from common-sense knowledge." In this form of knowledge, they see a source of inspiration in its own right for MAS-designs, supplementary to making sociological metaphors usable for DAI. In grounded theory, the same attitude is expressed, for example, in the suggestion that we make use of "in-vivo codes" (i.e., that we express what is to be analyzed in the language of the fields being studied [Glaser 1978, p. 70; Strauss 1987, p. 33]).¹⁸

The most obvious explanation for the similarities in research style between MAS and SI can be found in their common philosophical background of American pragmatism. Recent SI explicitly draws on these intellectual roots and can be traced back to Mead, Dewey, and Peirce. In American MAS research, the references to pragmatism are more implicit, but still evident, as in Gasser's (1991, p. 134) criticism of deductive logic as a basis for MAS processes: "[T]here are several problems with using deductive logic as a foundation for problem solving in open systems . . . since any deductive theory depends upon precursors such as . . . a model, . . . it doesn't seem unreasonable to say that when multiple viewpoints are at stake, logic may fail." This is reminiscent of Peirce's remarks in justifying abduction.

To summarize, we can explain how the research cooperation between DAI and current SI works by employing the analytical framework which has been developed in recent interactionist studies. In the process, the research context becomes the empirical case in which we can study the emergence of practical problem-solving in heterogeneous arenas of cross-disciplinary research. However, one should note that, to date, MAS and the related parts of SI have not converged to form a new hybrid scientific discipline. Or, to rephrase it in terms of social worlds theory, the two intersecting social worlds of SI and MAS have not formed a new social world such as "interactionist-MAS research".¹⁹ The actors continue to operate within the frame of reference of their respective disciplines. Quoting Malsch (1997), we might term this "computer reference" and "social reference." The research questions of common interest presented above represent an important link between the two intersecting

social worlds. While their meaning in computer sciences differs significantly from their meaning to sociologists (e.g., being relevant from a constructive perspective in one case, and from an analytical perspective in the other), they nonetheless focus attention in the same direction. Thus, we can regard them as “boundary objects” in the sense in which Star coined this term. More precisely, they represent the specific type of boundary object which Star and Griesemer (1989, p. 410) termed an “ideal type”: “It is abstracted from all domains, and may be fairly vague. However, it is adaptable to a local site precisely because it is fairly vague; it serves as a means of communicating and cooperating symbolically—a ‘good enough’ road map for all parties.” I would call this type of collaboration *trans-disciplinary* in order to distinguish it from the common picture of interdisciplinary as a mixture of disciplines.²⁰

SUMMARY AND DISCUSSION

The conceptual and organizational links between MAS research and SI have developed mostly in the U.S. and within a small community of researchers from both social and computer sciences. One cannot overlook the fact that in other countries, with their own scientific cultures, MAS research has found inspiration elsewhere and has focused on other aspects of MAS. Without going into detail, we can see that in European MAS research norm-oriented approaches, rational-choice theory, and systems theory have found more resonance in approaching the society-construction problem. However, looking at the importance of the results the connection between American DAI and SI has produced, this going concern can be seen as an exceedingly successful example of cooperation between sociology and technological sciences.

Concerning the question raised by Malsch, about the mode of migration of metaphors from the social sciences to DAI, the case discussed in this paper reveals that the notion of a “migration of metaphors” itself is questionable. As pointed out above, the contribution of SI to MAS research cannot be reduced to a couple of ideas transferred from one domain to the other. Instead, it is the mutual creation and maintenance of organizational structures in which common problems become solvable through two different branches of science pooling their resources—with both sides benefiting from the different disciplinary perspectives of the other. It is my belief that this enables us to understand the influence of SI on American MAS research. The seemingly parallel development of concepts in both social worlds turns out to be the result of successful trans-disciplinary collaboration based on both a shared pragmatist attitude in research style and research questions addressed by both disciplines, albeit from their own perspectives.

In discussing the cooperation between DAI and sociology, some authors stress the importance of the micro-macro issue (Malsch 1997; Florian 1998; Rammert 1998). When it comes to SI, however, some of these authors (e.g., Florian 1998) tend to “throw the baby out with the bathwater,” by suggesting that SI generally ignores structural aspects of society and that it has no concepts to deal with the micro-macro link. According to these authors, the specific problem of DAI is to cope with open networks consisting of not only heterogeneous but *a priori* unknown agents—rather than merely modeling systems like

closed groups of predetermined agents. Although this description of DAI's concerns is correct, we can still ask why an interactionist perspective should not be the appropriate way to approach these concerns. This criticism fails to realize the remarkable development current SI has undergone. Current SI, as described in this article, has undergone exactly the development which Joas (1988) claimed to be the way out of the "theoretical isolation" of earlier interactionist approaches. The small group of face-to-face actors has become an arena where representatives of different social worlds meet in order to deal with structural conditions through processes of negotiation (which at the same time reproduces *and* modifies these structural frameworks). The closed shop of the micro-social arrangement dissolves into an open, problem-oriented network, a going concern, consisting of all actors—even across sites—who prove to be sufficiently committed to the related problem. The apparently minor importance of the micro-macro link in current SI is mainly due to the fact that the latter avoids these less productive dichotomies right from the beginning (see Clarke and Gerson 1992, p. 180; Strauss 1993, p. 45 *passim*). The harsh confrontation of action and structure is replaced by a perspective which reveals the making of both in one dialectical process (see Gerson 1976, p. 796f).

The conceptual analogies observed in DAI and SI should not obscure the fact that all attempts to model socio-technical networks by drawing on sociological concepts have left important questions unanswered. As stressed by Rammert (1998), it is far from clear as to whether the models most suitable in conceptualizing social phenomena are equally appropriate in solving DAI problems. Thus, it would be shortsighted to regard the cooperation between DAI and sociology as simply a matter of translating social theory to machine systems. Rather, in computer science as in sociology, a shared concern and attitude is required toward both designing distributed socio-technical systems to be socially and technically appropriate, and to come to a suitable sociological analysis of these new phenomena. In this respect, the going concern of current SI and American DAI research provides us with an interesting example which in its contents and in its methods shows how to solve the practical and the analytical task in a consistent manner.

Acknowledgments: Helpful comments on earlier drafts of this paper were provided by Norman Denzin, Carl Hewitt, Werner Rammert, and Leigh Star. James Murphy and Leigh Star made it "readable" to an American audience. I am grateful to these colleagues. However, the remaining mistakes are mine.

NOTES

1. While "DAI" signifies the whole field of distributed computing, those researchers who focus on systems of partly autonomous, interacting agents in real-world settings are represented by the label "multi-agent systems" (MAS). Henceforth, I will concentrate on the latter.
2. "Agents" can be thought of as the technical counterpart of what we call "actors" in social reality. An agent is a component in a distributed software system which is able to take action largely independent from other parts of the system. It is important to note that agents are not completely pre-determined by system designers. Designing MAS does not consist of programming a complete set of agents; instead, it means to build a general system in which agents can evolve, act,

and change their properties. Nonetheless, in the end, all agents are a remote result of human programming work.

3. Some basic aspects of this concept originate in the mid-thirties; in 1957, it was fully developed (Hughes 1971, p. 52).
4. Current SI would not mind counting even technical or other non-human actors among the committed entities (see Star 1995, p. 15; Clarke 1995; Clark and Gerson 1992, p. 198; Bowker and Star 1996; Timmermans, this volume).
5. By reading Mead more carefully, however, it is obvious that he meant social *practice*. In *The Philosophy of the Act*, he stresses that "language is ultimately a form of behavior and it calls for the rationally organized society within which it can properly function" (1938, p. 518). In other words, speaking is social action and language is a medium for organizing society (see also Clarke 1990, p. 17).
6. Gasser's point is that there are quite a number of sociological concepts which are potentially valuable for MAS, but, as he demonstrates with Becker's (1960) commitment/side bets concept, these approaches often involve structures which current information science is not yet able to translate into program structures.
7. And it was never intended by Kornfeld and Hewitt's use of the scientific community metaphor. Indeed, they stress the importance of concurrency and negotiation on mutual control over resources in scientific communities. In addition, they further state that ascribed properties like pluralism, commutativity, or parallelism "are only incompletely achieved in practice" (1988, p. 312).
8. In answering the obvious question of why he did not choose another, more general problem-solving process as his metaphor, Hewitt, in his first article on the scientific community metaphor, notes that the latter is especially useful because scientific communities are highly formalized and standardized which allows for an easier transfer into technical contexts (Hewitt 1977, p. 350).
9. SI draws on Mead's social psychology, a conceptual system which—like Dewey's ([1896] 1965) reflex-arc critique—aimed to overcome behaviorist stimulus-response concepts of action. However, we should avoid a common misunderstanding: Mead's social psychology is not identical with SI, nor has it been completely integrated into traditional interactionist thought. It did, however, play an important role in the history of SI theory as it developed at the University of Chicago. In his programmatic work on a theoretical and methodological perspective for SI, Blumer (1969) explicitly built on Mead's concept of symbolically mediated action. However, his work was still a specific interpretation of certain parts of Mead's writings. If Blumer's perspective overemphasizes verbal face-to-face interaction and neglects larger structural conditions of society (to summarize standard criticisms of SI), then this is predominantly due to Blumer's historical situation. Writing in the 50s and 60s, Blumer aimed to establish a pragmatic antithesis to the reigning functionalist type of sociology in which people were described as nothing but executing agents of a dominating system of rules, norms, and other structures. In light of this confrontation with a deterministic type of sociology, it was no surprise that, in the first decades of SI, many researchers aimed to re-establish research on face-to-face interaction.
10. That is why these interactionists reacted with interest and cooperation rather than disapproval when Callon and Law (1982; Callon 1985) and Latour (1987) posed their thesis of a "symmetrical anthropology"—very different from Collins's harsh and dismissive reaction. Instead, interactionists and advocates of actor-network theory established a productive discourse resulting in some degree of approximation of both viewpoints (Star and Griesemer 1989; Fujimura 1992, 1991; Callon 1994; Bowker and Star 1996).
11. In the author's country, choosing a pragmatist perspective is not self-evident.
12. Speaking of differences between single machines and systems or networks of distributed artificial intelligence does not mean that we must regard DAI systems as consisting of spatially sep-

arated components. Instead, Davis and Smith (1983, p. 67) distinguish between spatial and functional distribution (i.e., even a physically integrated machine could logically consist of a countless crowd of agents—like those well known “wizzards” or “assistants” in application programs).

13. W.I. Thomas made similar statements in 1917; see Stone and Faberman 1970, p. 54ff.
14. The latter is easily observed by comparing the acknowledgments in articles and books by these researchers.
15. It is not only the dichotomy of social and technical but also the distinction between development and use which is challenged by interactionist thought (Bendifallah et al. 1988, p. 2).
16. In 1950, Alan Turing proposed a test setting in which an interrogator tries to find out which of the two other hidden actors (*A* and *B*) is male and which is female. *A* was meant to be replaced by a computer in order to find out if computers are able to think. The problem of this behavioristic test setting is that—far from being realistic—it works only within the entirely closed world of a game and its definite rules. For a profound discussion of this test and its epistemological consequences, see Heintz 1993a, p. 261.
17. I call this a research *style* because the approach Glaser and Strauss presented in 1967, and the refinements they have offered since its introduction, is not a strict methodology or merely a set of methods but a complete attitude toward empirical research, and it integrates organizational aspects of the research process (Glaser and Strauss 1967; Glaser 1978; Strauss 1987).
18. In light of recent results from STS studies, the heavy criticisms of common-sense knowledge as, for instance, a “naïve picture of the social” (which some sociologists state) becomes questionable. In my view, it is more appropriate to think of a continuum of enacted knowledge gradually differentiated by its suitability for being formalized.
19. For intersection and segmentation see Strauss (1978b, p. 122 and 1984). Intersection processes in the computing world are widespread; Kling and Gerson (1978) already discussed this phenomenon some 20 years ago. Currently, besides the DAI-SI link, there are other trans-disciplinary cases of cooperation to mention: for example, between ethnomethodology and computer sciences in the field of computer-supported cooperative work (CSCW; Suchman 1995) or between SI and computer sciences in the design of large scale systems (Star and Ruhleder 1996).
20. A good example for the divergence of perspective on mutually interesting concepts is provided by Gasser’s (1991, p. 113) assessment of commitment and similar interactionist concepts as “conceptually fruitful, though not presently computational.”

REFERENCES

- Adler, Mark R., Alvah B. Davis, Robert Wehmayer, and Ralph W. Worrest. 1989. “Conflict-Resolution Strategies for Nonhierarchical Distributed Agents.” Pp. 139–61 in *Distributed Artificial Intelligence*, Vol. II., edited by Michael Huhns and Les Gasser. London: Pitman.
- Agha, Gul. 1986. *Actors: A Model of Concurrent Computation in Distributed Systems*. Cambridge: MIT Press.
- Becker, Howard S. 1960. “Notes on the Concept of Commitment.” *American Journal of Sociology* 66: 32–40.
- Bendifallah, S., F. Blanchard, A. Cambrosio, J. Fujimura, L. Gasser, E. M. Gerson, A. Henderson, C. E. Hewitt, S. L. Star, W. Scacci, L. Suchman, and R. Trigg. 1988. “The Unnamable: A White Paper on Socio-Computational ‘Systems.’” Unpublished manuscript, Department of Computer Science, University of Southern California.
- Blumer, Herbert. 1969. *Symbolic Interactionism: Perspective and Method*. Englewood Cliffs, NJ: Prentice Hall.

- Bond, Alan H. 1991. "Commitments and Projects." Pp. 26–35 in *Proceedings of the 8th Conference of the Society for the Study of Artificial Intelligence and Simulation of Behaviour*, 16–19 April 1991, University of Leeds, edited by Luc Steels and Barbara Smith. London: Springer.
- Bond, Alan H. and Les Gasser. 1988. "An Analysis of Problems and Research in DAI." Pp. 3–35 in *Readings in Artificial Intelligence*, edited by Alan H. Bond and Les Gasser. San Mateo, CA: Morgan Kaufmann.
- Bowker, Geoffrey C. and Susan Leigh Star. 1996. "How Things (Actor-Net)Work: Classification, Magic and the Ubiquity of Standards." Unpublished manuscript, University of Illinois at Urbana-Champaign.
- Callon, Michel. 1985. "Some Elements of a Sociology of Translation. Domestication of the Scallops and the Fishermen of St. Briec Bay." Pp 196–229 in *Power, Action and Belief: A New Sociology of Knowledge?*, edited by John Law. London: Routledge.
- . 1994. "Four Models for the Dynamics of Science." Pp. 29–63 in *Handbook of Science and Technology Studies*, edited by Sheila Jasamoff. London: Sage.
- Callon, Michel and John Law. 1982. "On Interests and their Transformation: Enrollment and Counter-Enrollment." *Social Studies of Science* 12: 615–25.
- Clarke, Adele E. 1995. "Modernity, Postmodernity & Reproductive Processes ca. 1890–1990. Or, 'Mommy, where do cyborgs come from anyway?'" Pp. 139–55 in *The Cyborgs Handbook*, edited by Chris H. Gray, Heidi Figueroa-Sarriera, and Steven Mentor. New York: Routledge.
- Clarke, Adele E. 1991. "Social Worlds/Arenas Theory as Organizational Theory." Pp. 119–58 in *Social Organization and Social Process. Essays in Honour of Anselm Straus*, edited by David R. Maines. New York: Aldine de Gruyter.
- . 1990. "A Social Worlds Research Adventure. The Case of Reproductive Science." Pp. 15–42 in *Theories of Science in Society*, edited by Susan E. Cozzens and Thomas F. Gieryn. Bloomington: Indiana University Press.
- Clarke, Adele E. and Joan H. Fujimura. 1992. "What Tools? Which Jobs? Why Right?" Pp. 3–44 in *The Right Tools for the Job. At Work in Twentieth-Century Life Sciences*, edited by Adele E. Clarke and Joan H. Fujimura. Princeton, NJ: Princeton University Press.
- Clarke, Adele E. and Elihu M. Gerson. 1992. "Symbolic Interactionism in Social Studies of Science." Pp. 179–214 in *Symbolic Interactionism and Cultural Studies*, edited by Howard S. Becker and Michael McCall. Chicago: University of Chicago Press.
- Corbin, Juliet and Anselm L. Strauss. 1990. "Grounded Theory Research: Procedures, Canons, and Evaluative Criteria." *Zeitschrift für Soziologie* 19: 418–27.
- Davis, Randall and Reid G. Smith. 1983. "Negotiation as a Metaphor for Distributed Problem Solving." *Artificial Intelligence* 20: 63–109.
- Denzin, Norman K. 1992. *Symbolic Interactionism and Cultural Studies*. Oxford: Blackwell.
- Dewey, John. [1896] 1965. "*Philosophy, Psychology and Social Practice*," edited and with a foreword by Joseph Ratner. New York: Capricorn Books.
- Durfee, Edmund H., Victor R. Lesser, and Daniel D. Corkill, 1988. "Coherent Cooperation Among Communication Problem Solvers." Pp. 268–84 in *Readings in Artificial Intelligence*, edited by Alan H. Bond and Les Gasser. San Mateo, CA: Morgan Kaufmann.
- Ellis Carolyn and Michael G. Flaherty (eds.). 1992. *Investigating Subjectivity: Research on Lived Experience*. Newbury Park, CA: Sage.
- Fisher, Berenice and Anselm L. Strauss. 1978. "The Chicago Tradition: Thomas, Park, and Their Successors." *Symbolic Interaction* 1: 5–23.
- . "George Herbert Mead and the Chicago Tradition of Sociology." Part 1: *Symbolic Interaction* 2: 9–24; Part 2: *Symbolic Interaction* 2: 9–20.
- Florian, Michael. 1998. "Die Agentengesellschaft als sozialer Raum (A society of Agents as a Social Space)." Pp. 297–344 in *Sozionik: Soziologische Ansichten über Künstliche Sozialität*, edited by Thomas Malsch. Berlin: Sigma.

- Fujimura, Joan H. 1992. "Crafting Science: Standardized Packages, Boundary Objects and Translation." Pp. 168–211 in *Science as Practice and Culture*, edited by Andrew Pickering. Chicago: University of Chicago Press.
- Fujimura, Joan H. 1991. "On Methods, Ontologies, and Representation in the Sociology of Science: Where do We Stand?" Pp. 207–48 in *Social Organization and Social Process: Essays in Honour of Anselm Straus*, edited by David R. Maines. New York: Aldine de Gruyter.
- . 1988. "The Molecular Biological Bandwagon in Cancer Research: Where Social Worlds Meet." *Social Problems* 35: 261–83.
- . 1987. "Constructing 'Do-able' Problems in Cancer Research: Articulating Alignment." *Social Studies of Science* 17: 257–93.
- Galliers, Julia Rose. 1990. "The Positive Role of Conflict in Cooperative Multi-Agents Systems." Pp. 33–46 in *Decentralized AI*, edited by Yves Demazeau and Jean-Pierre Müller. North-Holland: Elsevier Scientific Publishers.
- Gasser, Les. 1992. "An Overview of DAI." Pp. 9–30 in *Distributed Artificial Intelligence: Theory and Practice*, edited by Nicholas M. Avouris and Les Gasser. Dordrecht: Kluwer.
- . 1991. "Social Conceptions of Knowledge and Action. DAI Foundations and Open Systems Semantics." *Artificial Intelligence* 47: 107–38.
- Gasser, Les, Nicholas F. Rouquette, Randall W. Hill, and John Lieb. 1989. "Representing and Using Knowledge in Distributed AI Systems." Pp. 55–78 in *Distributed Artificial Intelligence*, Vol. 2, edited by Michael Huhns and Les Gasser. London: Pitman.
- Gerson, Elihu M. 1983. "Work and Going Concerns: Some Implications of Hughes' Work. Unpublished paper, Tremont Research Institute, San Francisco.
- . 1976. "On 'Quality of Life.'" *American Sociological Review* 41:793–806.
- Gerson, Elihu M. and Susan Leigh Star. 1986. "Analyzing Due Process in the Workplace." *ACM Transactions on Office Information Systems* 4: 257–70.
- Glaser, Barney G. 1978: *Theoretical Sensitivity. Advances in the Methodology of Grounded Theory*. Mill Valley, CA: Sociology Press.
- Glaser, Barney G. and Anselm Straus. 1967. *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Chicago: Aldine.
- Hayes-Roth, F. 1980. "Towards a Framework for Distributed AI." *Sigart Newsletter* 73: 51pp.
- Heintz, Bettina. 1993a. "Die Herrschaft der Regel. Zur Grundlagengeschichte des Computers (The Rule of the Rule. On the Basic History of Computers)." Frankfurt a. M.: Campus Verlag.
- . 1993b. "Wissenschaft im Kontext. Neuere Entwicklungstendenzen in der Wissenschaftssoziologie (Science in Context: New Trends in the Sociology of Science)." *Kölner Zeitschrift für Soziologie und Sozialpsychologie* 45: 528–52.
- Henderson, Kathryn. 1995. "The Visual Culture of Engineering." Pp. 196–218 in *The Cultures of Computing*, edited by Susan Leigh Star. Oxford: Oxford University Press.
- Hewitt, Carl E. 1986. "Offices are Open Systems." *Transactions of the ACM* 4: 271–87.
- . 1985. "The Challenge of Open Systems." *Byte* 10: 223–42.
- . 1977. "Viewing Control Structures as Patterns of Passing Messages." *Artificial Intelligence* 8: 323–64.
- Hughes, Everett C. 1971. *The Sociological Eye*. Chicago: Aldine.
- Joas, Hans. 1988. "Symbolischer Interaktionismus: Von der Philosophie des Pragmatismus zu einer soziologischen Forschungstradition (Symbolic Interactionism: From Pragmatism to a Tradition of Sociological Research)." *Kölner Zeitschrift für Soziologie und Sozialpsychologie* 40: 417–46.
- Kling, Rob and Elihu M. Gerson. 1978. "Patterns of Segmentation and Intersection in the Computing World." *Symbolic Interaction* 1: 24–43.
- Kornfeld, William A. and Carl E. Hewitt. 1988. "The Scientific Community Metaphor." Pp. 311–20 in *Readings in Artificial Intelligence*, edited by Alan H. Bond and Les Gasser. San Mateo, CA: Morgan Kaufmann.

- Latour, Bruno. 1987. *Science in Action. How to Follow Scientists and Engineers Through Society*. Cambridge, MA: Harvard University Press.
- Maines, David R. and Jay Charlton. 1985. "The Negotiated Order Approach to the Study of Social Organization." *Studies in Symbolic Interaction* (Supplement 1): 271–308.
- Malsch, Thomas. 1997. "Die Provokation der 'Artificial Societies' (The Provocation of 'Artificial Societies')." *Zeitschrift für Soziologie* 26: 3–21.
- Martial, Frank von. 1990. "Interactions among Autonomous Planning Agents." Pp. 105–19 in *Decentralized AI*, edited by Yves Demazeau and Jean-Pierre Müller. North-Holland: Elsevier Science Publishers.
- Mead, George Herbert. 1934. *Mind, Self and Society*. Chicago: University of Chicago Press.
- . 1938. *The Philosophy of the Act*. Chicago: University of Chicago Press.
- Müller, Jürgen (ed.) 1993. "Verteilte Künstliche Intelligenz. Methoden und Anwendungen (Distributed Artificial Intelligence: Methods and Applications)." Mannheim: BI-Wissenschaftsverlag.
- Nwana, Hyacinth. 1996. "Software Agents. An Overview." *Knowledge Engineering Review*, 11: 1–40 (<http://www.cs.umbc.edu/agents/introduction/ao/>).
- Prus, Robert C. 1996. *Symbolic Interaction and Ethnographic Research: Intersubjectivity and the Study of Human Lived Experience*. Albany, NY: SUNY Press.
- Rammert, Werner. 1998. "Giddens und die Gesellschaft der Heizelmännchen: Zur Soziologie technischer Agenten und Systeme Verteilter Künstlicher Intelligenz. (Giddens and the Society of the Brownies: On a Sociology of Technical Agents and of Multi-Agent-systems)." Pp. 91–128 in *Sozionik: Soziologische Ansichten über Künstliche Sozialität*, edited by Thomas Malsch. Berlin: Sigma.
- . 1995. "Soziologische Zugänge zur Künstlichen Intelligenz. (Sociological Approaches to Artificial Intelligence)." Pp. 7–36 in *Soziologie und künstliche Intelligenz*, edited by Werner Rammert. Frankfurt a. M.: Campus.
- Schulz-Schaeffer, Ingo. 1993. "Das Forschungsfeld 'Verteilte Künstliche Intelligenz.' Ein Literaturbericht (The Research Domain of DAI. An annotated Bibliography)." Unpublished paper, TU Hamburg-Harburg.
- Shibutani, Tomatsu. 1955. "Reference Groups as Perspectives." *American Journal of Sociology* 60: 562–69.
- Snow, C. P. 1993. *The Two Cultures*. Cambridge: Cambridge University Press.
- Star, Susan Leigh. 1996. "Working Together: Symbolic Interactionism, Activity Theory, and Information Systems." Pp. 296–318 in *Cognition and Communication at Work*, edited by Yrjö Engeström und David Middleton. Cambridge: Cambridge University Press.
- . 1995. "Introduction." Pp. 1–35 in: *Ecologies of Knowledge*, edited by Susan Leigh Star, Albany, NY: SUNY Press.
- . 1991. "The Sociology of the Invisible. The Rising Key of Work in the Writings of Anselm Strauss." Pp. 265–83 in *Social Organization and Social Process: Essays in Honour of Anselm Strauss*, edited by David R. Maines. New York: Aldine de Gruyter.
- . 1989. "The Structure of Ill-structured Solutions: Boundary Objects and Heterogeneous Distributed Problem Solving." Pp. 37–54 in *Readings in Distributed Artificial Intelligence*, Vol. II., edited by Michael Huhns and Les Gasser. London: Pitman.
- Star, Susan Leigh and James R. Griesemer. 1989. "Institutional Ecology." *Social Studies of Science* 19: 387–420.
- Star, Susan Leigh and Karen Ruhleder. 1996. "Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces." *Information Systems Research* 7: 111–34.
- Star, Susan Leigh and Anselm L. Strauss. 1997. "Layers of Silence, Arenas of Voice: The Dialogues Between Visible and Invisible Work." *Computer-Supported Cooperative Work*, in press.
- Stone, Gregory P. and Harvey A. Faberman (eds). 1970. "Social Psychology Through Symbolic Interaction." Waltham, MA: Xerox Publishers.
- Strauss, Anselm L. 1993. *Continual Permutations of Action*. New York: Aldine de Gruyter.

- . 1988. "The Articulation of Project Work. An Organizational Process." *Sociological Quarterly* 29: 163–78.
- . 1987. *Qualitative Analysis for Social Scientists*. New York: Cambridge University Press.
- . 1985. "Work and the Division of Labor." *Sociological Quarterly* 26: 1–19.
- . 1984. "Social Worlds and their Segmentation Processes." *Studies in Symbolic Interaction* 5: 123–39.
- . 1978a. *Negotiations: Varieties, Contexts, Processes, and Social Order*. San Francisco: Jossey-Bass.
- Strauss, Anselm L. 1978b. "A Social World Perspective." *Studies in Symbolic Interaction* 1: 119–28.
- Strauss, Anselm L., Shizuko Fagerhaugh, Barbara Suczek, and Caroly Wiener. 1985. *The Social Organization of Medical Work*. Chicago: University of Chicago Press.
- Strauss, Anselm L., Leonhard Schatzmann, Danuta Ehrlich, Rue Bucher, and M. Sabshin. 1963. "The Hospital and its Negotiated Order." Pp. 147–69 in *The Hospital in Modern Society*, edited by Eliot Freidson. New York: Free Press of Glencoe.
- Strübing, Jörg. 1997. "Symbolischer Interaktionismus revisited. Konzepte für die Wissenschafts- und Technikforschung (Symbolic Interactionism Revisited: Concepts for Science and Technology Research)." *Zeitschrift für Soziologie*, 26: 368–86.
- Suchman, Lucy. 1995. "Making Work Visible." *Communications of the ACM* 38: 56–64.
- Thomas, William I. and Dorothy Swaine Thomas. 1928: *The Child in America*. New York: Alfred A. Knopf.