Social Knowledge and Social Action: Heterogeneity in Practice

Les Gasser
Computational Organization Design Lab
Institute of Safety and Systems Management
USC, Los Angeles, CA. 90089-0021 USA
gasser@usc.edu (213) 740.4046

Abstract

This paper is about some of the social aspects of knowledge and action relevant to thinking in AI, and in particular the basic experience of multiple perspectives and integrating different kinds of local knowledge. It discusses ways of rethinking a number of familiar concepts including facts, interaction, knowledge, and organization, raising questions about how well we can currently capture their social dimensions conceptually, representationally, and computationally. It suggests several approaches to developing more complete computational models of these phenomena.

1 Introduction

This paper is about some of the social aspects of knowledge and action relevant to thinking in AI, and I want to explain a bit about how and why I these ideas came to interest me so deeply. The crux of the matter is that in my own intellectual life, I have been continuously involved in more than one "field" at once. Two of the primary arenas in which I've found myself have been computing—specifically experimental computer science and AI—and sociology-specifically symbolic interactionism and studies of the integration of technology and work. The original idea I had was that these seemingly quite different arenas actually had a lot in common, and it would be interesting to try to articulate what I found fascinating in each to people who were more focused on the other—to be a boundary-spanner. For example, when I first seriously encountered computing in the mid 1970s, the problematic of large-scale software development was central, and it certainly seemed that large-scale software 1) was developed by large groups of people, (see, e.g., [Scacchi 84]) and 2) might have many of the structuring characteristics of large-scale organizations of other types, when software was seen as collections of processes (cf., [Durfee, *et ai* 87; Fox 79; Fox 81; Gasser, *et ai* 87]). Similarly, some AI researchers were beginning to deal with questions of scale and concurrency. They turned their attention to multiagent and distributed systems, ostensibly to cope with increasing problem complexity and to reflect what they saw as basic characteristics of the "real-world" [Bond and Gasser 88]. These movements within AI seemed to me to offer a very happy marriage with what I was reading about certain branches of sociology—most especially symbolic interactionism (the very name itself seemed just right) and later with "science studies" - the sociology of science and scientific knowledge.

In the midst of this ferment, 1 quickly found that while the discourses in computing/AI and in sociology had stimulating and encouraging similarities, there were certain underlying assumptions that made them, in practice, virtually incommensurate. And, in a very real way, my personal experiences in the several worlds as a boundary-spanner mirrored the incompatibilities in the world-views. This experience of living in several worlds at once, then, is both the primary motivation for and the main point of the present paper: to help explain what it means for several world-views not to jive, and to provide some new ways of seeing (looking into) some basic issues of AI and computing through glasses of incommensurate perspective—-to get at the pragmatics of deeply-experienced heterogeneity.

The basic experience of multiple perspectives, what might be called a fundamental point, is simply these two observations: 1) experiences vary over place and time, and 2) experiences in different places and times interact, necessitating some sort of integration.

Elsewhere, I have discussed several notions that derive from these observations [Gasser *86;* Gasser 91; Gasser 92]. They include the existence of a fundamental tension between local and global knowledge and action; certain results concerning the impossibility of common semantics or shared representations; the nature of and strength behind commitment as a binding force in inter-agent relations, and the causal direction of goals, social laws, rules, programs, etc. in shaping behavior. Taken together, it seems to me, the impacts of a perspective shift that first takes heterogeneity, not commonality, as given, and that second, observes the joint and interlocked nature of so much human and human-computer interaction, are great. How can this patterned, ongoing, joint, interlocked behavior be explained and designed, without recourse to commonality?

The outcomes of long-term inquiries into these issues seem interesting to me, because they help to address questions such as what are the differences between ma-

chine and human societies? How does this difference translate into explaining the problematics of how societies of machines and people coexist together? How do we as societies develop the categories of analysis that sometimes keep people and machines separate, and that sometimes treat them in common or related terms? To me, these sorts of questions are central to understanding the integration of computing (and of machines in general) with human activity. They critically influence the processes of construction and use of machines (e.g., specifications, explanations for errors and for proper behavior, etc.). In a sense, I think it is important to build bridges between the *social worlds* of system development, use and comprehension, and the *material worlds* of computational artifacts. I believe that these are not just exercises in formality, but are practical issues of high impact.

The overall aim of this paper, then, is to begin some new conceptual modeling, or in fact re-modeling, of some familiar concepts—to take some concepts familiar in AI such as communication, interaction, individual agency, knowledge, action, etc. and to recast them from a viewpoint that takes a heterogeneous social system as its basis.

Why would we want to do this? First, as an exercise in flexibility: trying to see things in new ways, and exploring alternative conceptions to those we are already familiar with. Second, because we are beginning to confront conceptual problems with current approaches when we approach the social aspects of agent behavior, as in the following circumstances:

- When we consider individual agents' cognition situated in group and social contexts [Rogoff and Lave 84].
- When we consider groups (e.g., organizations) as loci of action or knowledge (cf. IHutchins and Klausen 92; Weick and Roberts 92]).
- When we consider knowledge and activity that is distributed over space, time, semantics, etc. [Bond and Gasser 88; Gasser 91; Gasser 92]
- When we consider open versus closed systems questions, such as creativity and the generation of fundamentally new forms.

In order to examine this process of re-modeling, we need some objects of study.

2 Objects of Study for Social Analyses

I take one ongoing task of a mature science to be the specification of its objects of study—the delineation of the conceptual entities that are of interest for its scientists. What are the objects of study that raise questions of social knowledge and social action useful for progress and theory in AI?[1] I'd like to make some observations about what things or phenomena we might examine and seek to describe, analyze, and explain in computational terms, and to investigate what are some questions to raise about these things. This is a suggestive, rather than an exhaustive list; in particular it will leave out *individuals ' knowledge and behavior, social cognition,* that is, knowledge about social and organizational entities [Higgins, *et al.* 85]; *risks, costs and moral order,* and *goals or intention,* which I intend to treat at greater length elsewhere.

2.1 Facts and Social Facts

There are several kinds of alternatives to the relatively conventional (in AI) conceptualizations view that facts are statements about the world known to be true in all contexts (e.g., possible worlds semantics of knowledge and belief). These include 1) the treatment of facts as continuously reinterpreted statements with dynamic "facticity"-statements repeatedly transformed, reinforced, and re-valued as they are incorporated in ongoing discourses (cf. [Latour and Woolgar 79]) so that their stable or reified character is a *product* of action, rather than simply a basis of action, and 2) Durkheim's notion of "social facts." Durkheim described social facts as "ways of acting, thinking, and feeling" that exist outside individual consciousness, that are diffused widely within a group, and that exert "a coercive power" over the activities of individuals, "recognizable by the resistance that it offers any individual action that would violate it." He points out that when taking on certain social commitments, "I perform obligations which are defined outside myself and my actions...we are ignorant of the details of the obligations we must assume, and ...to know them we have to consult the legal code and its authorized interpreters...the above statements will apply [to] each member of a society in turn" [Thompson 85], pp. 68-71. The point is that social facts reside in collectivities, not in individuals: "The determining cause of a social fact must be sought among antecedent social facts, and not among states of individual consciousness" [Thompson 85], pp. 86].

2.2 Individuals (Facts, Agents, etc.)

Individual knowledge, performance and achievement has long been the focus of AI (cf. [Bond and Gasser 88; Bobrow 91]). But much has been taken for granted. What is the nature of the individual agent? In what sense is it possible to conceive of an individual, carving one out of a continuous web of social interaction and involvement? Said another way, what aspects of individuals are *not* social facts? What is the boundary of any individual, in terms of action, time, knowledge, perception, etc., and how is our knowledge of these boundaries constituted *non-socially?*[2] How do stable individuals emerge in the collective action of societies and organizations? .

[1] In describing the study objects that follow here, I mean the words "individual," "organization,* etc. to refer to either people or machines; I mean this as a provocative stance.

[2] Once you think you have a clear answer to the boundary question, consider individuals as aggregates of parts—as (de)composable systems—and see if your answer holds up! See below...

[3] For the uninitiated, the notion of individuals emerging in collective activity may seem strange. Two examples: stable software processes built and maintained by software teams

Gerson has presented a simple and cogent conception of individual as "something for which nothing else will substitute for *each* and *every* purpose" [Gerson 91] pg 1. He points out that any conception of an individual thing depends on a *recognizer* (who assesses the substitutability and differentiation of the individual thing), and that the ongoing process of recognition is subject to mistakes. We discover and correct these mistakes, in general, due to the restrictions on action that they entail. Suppose a medical-diagnosis knowledgebase is mistakenly loaded into a circuit-diagnosis system, and doesn't substitute. The diagnosis system does *somethin-g*—maybe it beeps, crashes, or emits an error message—but it doesn't cooperate with its user in diagnosing circuits. Trying to treat a painter like a car mechanic won't work, because we depend on the painter's participation in fixing the car, and it's not forthcoming. In Gerson's words, "in specific local circumstances we live in a world of alliances which corrects mis-identifications" [Gerson 91] pg. 1 (cf. Durkheim's concept of the coercive power of social facts, mentioned above).

Thus, says Gerson, non-substitutability "is a function of cooperation and response from others; *there is no single thing* in general...We can reliably recognize something as *an* individual and as the same individual only if there is equivalence of criteria across recognizers [over time and place]. This is achievable only for very narrow purposes and for relatively short periods of time." Moreover, "some things [e.g., people] can actively manipulate the process of recognizing [by how they] anticipate and negotiate the criteria which others use to recognize them. They can insist on some criteria and rule out others....they can decide to be another individual, or [to be] individual in another way. When this happens we have things recognizing or constructing each other as individuals, the identity of each being dependent upon its cooperation with the identity of the other. In this situation, things demand recognition of their identities on their own terms as the price of cooperation" [Gerson 91] pgs. 2-4.

2.3 Interactions

What is the nature of interaction among individuals? Do we need a clear and delineated conception of the individual in order to conceive of interaction? For example, once we have located the very nature of individual agenthood in social processes—once agents become social facts—against what ground are we to give semantics to messages which travel across time and place between agents? What are the boundaries of interaction? For instance, suppose a sending agent gives notice to a receiving agent that "a proposal will arrive in a following message." Where and when does the interpretation of the proposal message begin and end? Does it begin with the notice message? Does it begin with prior messages, activities, and world states that over time generated the internal and external structures that allowed for interpretation of the notice message and subsequent assimilation of the proposal message? How do we separate the interpretation of a message from the activity and structures that establish the context in which it is interpreted? (cf. Gerson's note on the ways agents can influence their own substitutability and identity). It is certainly possible to set up very complex interpretation structures beforehand and to reduce interaction to sending a very small set of tokens, or even to sending none [Genesereth, *et al.* 84]—that is, to place more and more of the import of the message into what Kenneth Gergen has termed the retrospective and emergent contexts of action (see below). It would seem that the nature of the boundaries of interpretation, hence of the meaning of "message" is contingent on the socially-emergent definition of "message" and the nature of these contexts.

Of course, the same is true of other kinds of interactions. Gergen has indicated the difficulty of interpreting and even of identifying social action, using the following example:

> "If I see my friends Ross and Laura approach each other at a social gathering, and Ross reaches out and momentarily touches Laura's hair, precisely what have I observed? What action has occurred before me? How am I to identify it?" [Gergen 82] pg. 60.

Gergen points out that while precise measurement of the movements (e.g., assessments of physical state change) might be possible, they would tell us little about what the movements mean in terms of ongoing relations at the social gathering, before it, and after it. That is, any such measurements would tell us little about how to act in relation to Laura, Ross, and others. He points out that (and gives long examples of how) both "retrospective context" (things that have occurred and interpretations that have been made before the observed action) and "emergent context" (occurring after the action) play key roles in our interpretations. As we piece together these contexts, new information may trigger arbitrary revisions of prior interpretations. (We should note, following Durkheim, that all of these interpretations and the processes of their revision are subject to the coercive power of social facts.) This leads Gergen to three conclusions (quoted from [Gergen 82] pgs. 62-63; [Collins 90] and [Baker and Hacker 85] make many similar points):

1. "The identification of any given action is subject to infinite (post-hoc) revision."
2. "The anchor point for any given identification is not fundamentally empirical, but relies on a network of [other] interdependent and continuously modifiable interpretations."
3. "Any given action is subject to multiple identifications, the relative superiority of which is problematic."

The basic point of conclusion following Gergen's analysis is that actions by themselves (that is, measured changes of state, the most common representations of action in AI) are next to useless as a ground for identifying

are very clearly individuals that continuously (re)emerge in social processes. Similarly, people are products of collective action in very physical ways—food, clothing, shelter, health care, etc. are all continuously and collectively (re)arranged, and the knowledge involved in these activities is no less so.

and understanding social activity. This is especially the case since so much social activity is a matter of *symbolic interaction,* not physical constraint, and such symbolizations are Durkheimian social facts, situated with ongoing activities. Where is identification and specification of action (and, by extension, of state) to be "grounded?" The perception, identification and interpretation of action from a social standpoint is a matter of classification within a system of meanings, that is, within a localized classification system which itself is a collection of social facts (and thus has socially coercive power). As W.I. Thomas has pointed out, "Things perceived as real"—that is, perceived through a socially-factual system of classification and interpretation—"are real in their consequences" [Thomas 23]. What we haven't yet been able to accomplish squarely, it seems to me, is to model computationally the dynamic, social, and multilayered nature of classification and interpretation processes, so that they are intimately coupled with (and become) a coercive social power. The stability of our networks of interaction comes, in general, by predefinition, through predefined, regular computational structures, rather than by emergence and social interaction among participants.

2.4 Pragmatic Knowledge

Suppose we consider the pragmatic knowledge of a system to be the ability of that system to produce certain patterns of outputs when given certain patterns of inputs. For example, if a particular knowledge-based design system is given information about technical features of a manufacturing technology, organizational goals, etc., and manages to produce acceptable human infrastructure designs [Gasser *et ai* 93]), we could say, intuitively, that the system "knows how" to produce useful designs. That is, it produces input-output mappings judged effective by some evaluator. If a vehicle tracking system is given a set of low-level sensor data and produces accurate interpretations of that data as vehicle tracks, we will say that the the analysis system "knows how" to produce those interpretations (cf. [Durfee, *et ai* 87]. This view of the knowledge of a system locates that knowledge in what the system *actually does in context*—not in what is the lasting truth-value of statements in its knowledgebase[4]. That is, from this vantage point, the pragmatic knowledge of a system encompasses all the circumstances of that system's use, including the process by which its results are judged to be acceptable, the resources it uses, and even (in the extreme) that fact that the system was supplied with power—for without these things, by our definition, the system would not "know how" to produce interpretations or designs.

With this pragmatic view of knowledge as "knowing how to do," we ground knowledge in the practical behavior and context of a system. If the system has "all the facts," necessary to address a problem (and so is "epistemologically adequate," in McCarthy's terms), but cannot move deductions based on those facts from one place to another within its internal organization, then the system *doesn't* "know how" to produce its results. If a result is described by a quality or time constraint, and the system cannot structure itself to meet that constraint, then, the system does *not* "know how" to produce its results[5]. Moreover, if the interpretation of a system's output changes after the output has been produced (emergent context), then that system's knowledge has also changed. Thus pragmatic knowledge is not solely located "within" the system, but is the outcome of the engagement of the system with its inputs, and with its interpretive (i.e. retrospective/emergent) and resource contexts.

One way to see part of this notion computationally is to think of flexible distributed production systems such as the one discused in [Ishida, *et ai* 92]. This system allowed flexible reconfiguration of production system-based agents (their number, identity, and knowledge boundaries) in response to changing environmental demands and resources. When faced with time, quality, or resource demands it could not meet, the system reorganized its agent-to-knowledgebase associations (including facts, rules, and agent models), adding or deleting agents until the demands could be met or until the overall system resources were overtaxed. That is, the pragmatic knowledge of the system—being able to produce results that met external criteria—was a matter of the *dynamic system organization tn the particular environmental context,* and not simply of the facts, rules, and mapping knowledge in the agents' databases. Changing the environmental demand meant that knowhow disappeared and then was recovered as the system adapted its structure.

2.5 Organizational Cognition

How can concepts such as action, cognition, perception, and memory be conceived where organizations and groups are the active, cognizing, perceiving, remembering entities (cf. [Hutchins and Klausen 92; Weick and Roberts 92])? For each of these concepts, where is its locus, and what gives it its stability or pattern? How does organization emerge along with the collective action of individuals?

First, let us consider the issue of aggregation: how to "put together" collections such as knowledge-based processes, "agents," and/or people into an organized whole, and how to have them act together in response to some higher-order phenomenon—that is, a phenomenon at the level of the whole, not at the level of the components. What would this look like?

We can think of aggregation as having the following four aspects. First, there would be some identifiable entities that are put together. Second, these entities would be individually responsive to some environmental circumstances on their own—there must be a way of

[4]To me, this makes more sense anyhow, because facts (statements) in and of themselves are useless; any "fact," to be useful, must be applied, in a context—that is, it must somehow be incorporated into action. This is not to say that a system with such know-how is autonomous in its knowledge, or that somehow the knowledge is solely embedded in the system itself—in fact my argument is just the opposite.

[5]This is not entirely like McCarthy's heuristic adequacy, because it encompasses factors outside the system.

talking about them as individuals, with respect to some class of environmental stimuli and substitutability criteria. Third, there would be some mechanism or process that welds them together into an ongoing unit that exhibits some routineness, stability, or pattern. Fourth, this higher-order unit would itself respond in some patterned way to some qualitatively different class of stimuli, such that the overall response of the aggregate is different from the response of the individual units. That is, the group of individuals will not be substitutable for the aggregate with respect to the ongoing aggregate-level environment, and the responses and character of the individuals would be different by virtue of their participation in the organization.

From this description we can see there is some relationship between the interconnecting process and the class of higher-order stimulus that defines the nature of the aggregation. We can also see that the defining characteristic of an aggregate is 1) that it is a higher-order patterned response, which means that the interactions among parts must also be patterned to some degree, and 2) that the identifying character of the aggregate is determined in part by the character (and level) of the stimulus and response; in effect, the environment has a hand in defining whether something is an aggregate or not. (This is in line with the previous discussion of individuality as non-substitutability.)

Note that nowhere have we spoken of the members of the aggregate having any sort of (common) goals or intentions. To be identified as an aggregate vis-a-vis some observer or interactor, it is sufficient that there is an overall pattern to the members' collective activity in response to a class of stimuli, and we need not attribute any notions of "cooperation" or "working together" to that pattern[6].

Overall, for thinking about socially-constituted knowledge and action, we would like to avoid the notions of goal and intention because we want to deal with multi-level aggregates at multiple and arbitrary levels of aggregation. In such structures, concepts such as goals and intention become problematic, because we don't have a clear idea of where to situate responsibilities (e.g. of parties for goals, when parties are aggregates) or how to allocate action (e.g., for achieving goals) when action is distributed and simultaneous.

A number of proposals and approaches to aggregation exist for AI systems. These include:

- Aggregation using commitment and action restrictions [Levesque, *et al* 90].
- Aggregations using federation, that is, coupling through standard languages, or interaction interfaces, such as those proposed for Enterprise Integration and knowledge interchange [EIF 92].
- Aggregation using metaobjects and metalevel information, including representations or models of other agents for prediction and interpretation as used commonly in DAI approaches to coordination [Durfee, *et al* 87; Gasser, *et al* 87; Gasser and Briot 92]

As currently described, all of these still require a-priori common objects, namely the interaction languages, metaconcepts, and behavioral rules or programs that will cause agents to conforms to the standards. How can we think of an alternative?

What we need to look at is not shared metalanguages for interaction, but instead the processes through which new metalanguages become generated, accepted and integrated. A particular interaction standard is not the interesting part—the interesting part is the standard formation process itself; this is what we need a model of. This is because, first, we should begin to think about aggregation as not as state but instead as an ongoing process of aggregating (cf. [Weick 79]). A consequence of this is that boundaries of aggregates are always fluid, in terms of knowledge and action. Second, in successful precesses of aggregation, something keeps the parts together over time. That is, something keeps the aggregating process "on track." Aggregation is then not a state, it is a process of keeping things on track. How is this accomplished? We might suggest that if can be accomplished in several ways:

- By setting up social barriers to participation (social coercive power): if a participant doesn't conform to the recognized interaction patterns, it doesn't get into the game. How do these patterns come about? One way to think about this is that those agents who already have a game going in some way that get to set up new games with new requirements. One would expect that in open systems, aggregations would build around existing islands of compatible behaviors, rather than in a completely ad-hoc fashion (and this reflects observed patterns in biology and in social groups).
- By establishing open systems of checks and balances for assessment and social control. Rather than hierarchical systems, in which distributed control meets singularities at the top and bottom of the hierarchy, we need to explore non-hierarchical control regimes in which agents control one another via distributed checks and balances.
- By balancing degrees of reification and control with environmental pressures for disaggregation—that is, intentionally allowing interpretations and actions to vary for different participants, and allowing room for local adaptation, e.g. by activities of fitting, augmenting, and working around [Gasser 86]

  For example, in the context of a theory-based organization design/analysis system under construction at USC [Gasser *et al* 93], we have found that system interactions tend to be differentially interpreted in different organizational and social contexts, yet those contexts must interrelate. Thus the same system inputs and outputs must serve as bridges across organizational cultures. This has meant I) a conscious decision to use loosely-defined user input categories and qualitative input measurements (such as high-medium-low), so that the definitions

[6]This also opens the door to reconceptualizing goals and intentions as simply repeated patterns of action, rather than as mental states; details will be left to a future paper.

more easily fit into local users' interpretations, 2) user-tailorability of the definitions of the system's key conceptual vocabulary, and 3) a model of design as alignment between user-interpreted criteria. Thus the social use of a system fosters reconfiguration of its meanings. (Additional examples of this sort of reconfiguration in the context of error redefinition and system workarounds are provided in [Gasser 86].)

3 Conclusions and Some New Research Actions

Trying to use individual programs as a basis of action in distributed, situated, or group-based systems is like trying to use individual facts as a basis of action: they are insufficient in themselves. Thus, the idea of building a distributed community out of a pre-existing collection of programs is inadequate, when we treat programs as facts—programs must be evolved by the community at all levels. Of course, in reality this is what happens, if we consider the actual trajectory of any program, and not just at its autonomous behavior. What makes it work in practice is the open human activity in the marketplace of its use and evolution.

In a similar vein, the notion of substitutability in flexible social contexts is key to understanding knowledge interchange and interoperability of heterogeneous systems. Using programs in different contexts is a social process, and can't be located in the interaction languages or individual programs themselves.

Most DAI research works from the premise that some stable set of agents with stable architectural boundaries come together and coordinate their activities in the solution of joint problems. That is, a stable society of agents emerges from the constructive interactions of multiple pre-existing members. The primary problem, then, is how to design the individuals so that they can effectively coordinate when enlisted in joint problem-solving situations. In the social theories underlying these systems, social roles and social-level effects are founded in individual action and knowledge.

The argument in this paper is *not* that machines and mechanized knowledge or action necessarily are and must be inherently a-social. Instead, I have been trying to show, first, something of what we might need to do to incorporate greater sociability into machines, and second, some more directly social angles for thinking about the machine/human ensembles that we do work with. Unlike the stances of Collins, Dreyfus, Searle, etc. [Collins 90; Dreyfus 79; Searle 84], I'm not saying that for AI to be complete in some sense, artificially intelligent agents must be fully socialized. Instead, I'd like to suggest that with computers as partners, we have several opportunities to explore alternative theoretical models of sociability and culture, namely, the varieties of society and culture that emerge among collections of semi-autonomous machines and people-machine ensembles. I suggest seriously treating these as alternative, model cultures and societies, to learn more about how far our current conceptualizations of culture and society go.

Here are some suggestions for how to take some concrete steps toward more fully social yet nonetheless computational models:

- Learn to describe, build, and experiment with communities of programs 1) that generate, modify, and codify their own local languages of interaction; 2) in which kinds and degrees of structure and reification and both increase and decrease with use; and 3) that modify both their knowledge and their activity structures at all levels of analysis—i.e., communities of programs that evolve the languages in which they are written.
- Define and demonstrate social mechanisms of dynamic category formulation, classification, and concomitant reification—the active (re)formulation of agreed-upon basic concepts and their use in joint interpretation and discourse processes. "Social mechanisms" would be those in which categories, classification activities, reifications, structures, etc. were subject to Durkheim's social coercion.
- Investigate how these and similar reasoning, modeling, explanation, and activity structures—e.g., dynamic aggregation, reification, etc.— appear in other fields and contexts, including biology (evolutionary, developmental, and theoretical biology), formulation of scientific knowledge, and social control/social change processes.

4 Acknowledgements

Both the style and substance of this paper have been significantly influenced by the work of Leigh Star, who, with many others, has eloquently argued the necessity of making work and responsibility visible. The other primary influences on my thinking here have been Phil Agre and, as usual, Elih Gerson. The actual writing was simplified greatly by the generous encouragement of Leslie B. Hill.

References

[Baker and Hacker 85] G.P. Baker and P.M.S. Hacker, *An Analytical Commentary on the Philosophical Investigations,* Blackwell, Oxford, 1985.

[Bobrow 91] Daniel G. Bobrow, "Dimensions of Interaction," *AI Magazine,* Fall, 1991.

[Bond and Gasser 88] Alan H. Bond and Les Gasser, "An Analysis of Problems and Research in Distributed Artificial Intelligence" in Alan H. Bond and Les Gasser (eds.), *Readings in Distributed Artificial Intelligence,* Morgan Kaufmann Publishers, San Mateo, CA, 1988.

[Collins 90] Harry M. Collins, *Artificial Experts: Social Knowledge and Intelligent Machines,* MIT Press, Cambridge, MA., 1990.

[Dreyfus 79] Hubert Dreyfus, *What Computers Can't Do,* Harper and Row, New York, 1979.

[Durfee, *et al* 87] Edward M. Durfee, Victor R. Lesser, and Daniel D. Corkill, "Coherent Cooperation among Communicating Problem Solvers," *IEEE Transactions on Computers,* C-36, September, 1987.

[EIF 92] Proceedings of the First International Conference on Enterprise Integration, Hilton Head, South Carolina, MIT Press, June, 1992.

[Fox 79] Mark S. Fox, "Organization Structuring: Designing Large, Complex Software," Technical Report CMU-CS-79-155, Computer Science Department, Carnegie Mellon University, December, 1979.

[Fox 81] Mark S. Fox, "An Organizational View of Distributed Systems," IEEE Trans. Systems Man and Cybernetics, SMC-11, pp 70-80, January, 1981.

[Gasser 86] Les Gasser, "The Integration of Computing and Routine Work," ACM Transactions on Office Information Systems, 4:3, July 1986, pp. 225-250.

[Gasser, et ai 87] Les Gasser, Carl Braganza, and Nava Herman, "MACE: A Flexible Testbed for DAI Research," in M.N. Huhns (ed.), Distributed Artificial Intelligence, Pitman Publishers, London, 1987.

[Gasser 9I] Les Gasser, "Social Conceptions of Knowledge and Action," Artificial Intelligence, 47, January/February, 1991.

[Gasser 92] Les Gasser, "Boundaries, Aggregation, and Identity: Plurality Issues in Multi-Agent Systems," in E. Werner and Y. Demazeau, eds., Decentralized AI S, North Holland, New York, 1992.

[Gasser et al. 93] Les Gasser, Ingemar Hulthage, Brian Leverich, Jon Lieb and Ann Majchrzak, "Computational Organization Design", COD Research Memo 9, Computational Organization Design Lab, Institute for Safety and Systems Management, USC, 1993.

[Gasser and Briot 92] Les Gasser and Jean-Pierre Briot, "Object-Based Concurrent Programming and DAI" in Nikos Avouris and Les Gasser (eds.), Distributed AI: Theory and Praxis, Kluwer, 1992.

[Genesereth, et al. 84] Michael Genesereth, Matthew Ginsberg, and Jeffrey S. Rosenschein, "Cooperation Without Communication" Stanford University Computer Science Department Technical Report HPP-84-36, September, 1984.

[Gergen 82] Kenneth J. Gergen, Toward Transformation in Social Knowledge, New York, Springer Verlag, 1982.

[Gerson 91] Elihu M. Gerson, "Individuals," Paper presented at the AAAI Spring Symposium on Composite Systems Design, Stanford, CA, 27 March 1991. Available from Tremont Research Institute, 458 29th Street, San Francisco, CA. 94131 (tremont@ucsfvm.ucsf.edu).

[Higgins, et ai 85] E. Tory Higgins, Diane N. Ruble, and Willard W. Hartup (eds.), Social Cognition and Social Development, Cambridge University Press, New York, 1985.

[Hutchins and Klausen 92] E. Hutchins and T. Klausen, "Distributed Cognition in an Airline Cockpit," in D. Middleton and Y. Engestrom, eds., Communication and Cognition at Work, Cambridge University Press, 1992.

[Ishida, et ai 92] Toru Ishida, Les Gasser, and Makoto Yokoo, "Organization Self-Design of Distributed Production Systems," IEEE Transactions on Data and Knowledge Engineering, 4:2, April, 1992,

[Latour and Woolgar 79] Bruno Latour and Steve Woolgar, Laboratory Life: The Social Construction of Scientific Facts, Sage, Beverly Hills, CA., 1979.

[Levesque, et ai 90] Hector J. Levesque, Philip R. Cohen, and Jose H.T. Nunes, "On Acting Together," in Proceedings of the Eighth National Conference on AI, pp. 94-99, July, 1990,

[Rogoff and Lave 84] B. Rogoff and J. Lave (eds.) Everyday Cognition: Its Development in Social Context, Harvard University Press, Cambridge, MA, 1984.

[Scacchi 84] Walt Scacchi, "Managing Software Engineering Projects: A Social Analysis," IEEE Transactions on Software Engineering, SE-10:1, pp. 45-59, January, 1984

[Searle 84] John Searle, Minds, Brains, and Science, Harvard University Press, Cambridge, MA., 1984.

[Thomas 23] William I. Thomas, The Unadjusted Girl, Boston: Little Brown, 1923.

[Thompson 85] Kenneth Thompson, Readings from Emile Durkheim, New York: Tavistock Publications Ltd., 1985.

[Weick and Roberts 92] Karl E. Weick and Karlene H. Roberts, "Organization Mind and Organizational Reliability: The Case of Flight Operations on a Carrier Deck," working paper presented at the Workshop on Adaptive Processes and Organizations: Models and Data, Santa Fe Institute, February, 1992.

[Weick 79] Karl E. Weick, The Social Psychology of Organizing (2nd ed.), Addison Wesley Publishing Company, Reading, MA., 1979.