

Generalized Media of Interaction and Inter-Agent Coordination

Thomas Malsch, Ingo Schulz-Schaeffer

Arbeitsbereich Technikbewertung und Technikgestaltung
Technische Universität Hamburg-Harburg
Schloßmühlendamm 32, D-21073 Hamburg
malsch@tu-harburg.de, schulz-schaeffer@tu-harburg.de

Abstract

This paper investigates social intelligence from a sociological perspective. Starting from social interaction and the problem of "double contingency" it shows how "generalized symbolic media of interaction" facilitate coordination processes in modern human societies. It advocates the concept of media as a source of fruitful inspiration for the design of socially intelligent agents and artificial social systems.

Introduction

In AI literature on social agents we are confronted with a confusing variety of concepts and approaches ranging from artificial insects to human-like agents, from software agents exchanging messages to robots interacting with each other via embodied sensory and motor skills, from friendly or benevolent agents to selfish or even aggressively competing agents. On a closer look, however, we can roughly distinguish between three different approaches to agenthood: the reactive agent, the embodied agent, and the cognitive agent. Each of these concepts of agenthood or agency is (implicitly or explicitly) related to a specific model of sociality or social coordination: the reactive agent to insect states, the embodied agent to ape communities, and the cognitive agent to communities of human experts. No doubt, these different approaches all have their particular merits and advantages, but each of them fails when it comes to explaining and modeling specific features of human societies.

Modern human societies pose a particular challenge to Distributed Artificial Intelligence (DAI), not only because they are much more complex than animal societies - they are also able to achieve coordination in spite of the fact that human beings are capable of unpredictable, voluntaristic action. The insect state metaphor is inadequate because coordination is based on instinctive behavior involving no degrees of freedom; benevolence is thus triggered by olfactory signals, whilst rudimentary task differentiation (as in honeybee colonies between cleaners, nurses, harvesters, gate-keepers) follows predefined biological patterns (Drogoul and Ferber 1994). The concept of embodied agency can show how solidarity and group cohesion in gorilla families is established and maintained by grooming and social rubbing, but it fails when it comes to explaining

how coordination is possible in large human societies with anonymous relationships, where direct face-to-face contact between members is but one mechanism among many for achieving coordination. Similarly, the cognitive agent approach, based on the idea of furnishing autonomous intelligent agents with models of other agents ("acquaintances"), cannot describe and explain coordinated action beyond the narrow scope of personally acquainted experts. In short, such approaches to agency and sociality are unable to elucidate patterns of "impersonal" social coordination, as they remain unaware of the distinction between community (characterized by direct communication, physical co-presence, and personal acquaintance) and society (characterized by functional differentiation, structural complexity, and generalized media of interaction).

In our paper we wish to discuss the problem of coordination from a sociological perspective. Let us assume that agents are "black boxes" operating at "arm's length relationships" (Hewitt 1991), i.e. agents who are unable to look into each others heads and unable to know about each others internal operations or mental states in terms of beliefs, desires, or intentions, but who are capable of observing one another's behavior: how can they coordinate their actions without knowing what the other agent wants and whether he (she/it) is hostile or friendly, selfish or benevolent? The answer is: by building up complementary expectations and stabilizing them in the form of so called "generalized media of interaction" (Parsons 1969).

Agency, Coordination and the Problem of Double Contingency

Building up successful coordination between agents interacting at arm's length distance is a fundamental design problem both in DAI and sociology. Taking agency as their point of departure, many researchers concerned with artificial social agents are far too occupied investigating what an agent is, what he wants and how he is able to attain his goals. However, we do not need to begin with an elaborate conception of agent properties. Starting with a weak notion of agency we only have to presuppose (1) that agents are able to act or behave in different ways and; (2) that they are able to select from their possible behaviors. Agents are

characterized by a certain degree of autonomy for as Anthony Giddens points out: "Agency concerns events of which an individual is the perpetrator, in the sense that the individual could, at any phase in a given sequence of conduct, have acted differently." (Giddens 1984: 9) In other words: agents are entities, whose behavior is contingent on their own selections among available alternatives and not directly controlled by external intervention.

The situation gets a little bit more complicated, however, when we move along from individual to social agents, that is to agents that do interact with each other. Social action can be understood as a specific form of situated action. Situated action is when an agent has to take into account the specific circumstances of a situation in order to perform his actions, because these circumstances impinge on his actions as enabling or constraining factors. Social action is situated action where these factors are the actions of other agents. Let us now assume that the agents involved in such a social situation are in the same position we started with: the only thing they know about each other is that the other entity is an agent, but they perceive each other as mere black boxes so that they can't be sure what the other agent will do next. In this context the problem of coordination between agents can be described basically as a problem of double contingency. In an interaction each agent's actions are contingent on his own selections among certain possible actions, but as reactions to the actions of other agents his selections are at the same time contingent on the selections of these other agents (Parsons and Shils 1951: 16). So if each agent is able to act in different ways and no agent knows for sure what the other agents will do next, how then is mutual adjustment of actions and coordinated interaction between agents possible?

Sociological research into human societies suggests two general answers to this question. The first answer is that double contingency in the pure form just described never actually occurs in the social world we live in. Each social situation is structured by cultural patterns that allow us to have reliable expectations about the actions of the other participants patterns that can be used as media of successful interaction and coordination. The second answer emphasizes the point that social situations are rarely completely structured by cultural patterns so that double contingency is indeed a real problem, a problem that has to be solved by generating complementary expectations in the process of interaction itself. The two answers are neither contradictory nor mutually exclusive. They both assert that patterns of social interaction are not constructed by simply adding or aggregating individual agents' programs of action but that they are rather an emergent result of generating complementary expectations. While the first answer focuses on expectations that have been developed, structured and stabilized over the course of history and that are now available as a common coordination resource, the second answer focuses more on the social process of shaping and forming reciprocal or mutual expectations in the course of a particular sequence of interactions. In this paper we are

mainly concerned with the first response to the problem of double contingency.

Complementary Expectations and the Generalized Media of Interaction

Without a common language no interaction between human actors would be possible. Whenever we utter a symbolic expression we presuppose that the other actor will understand it. However, given that the actors can't look into one another's heads how can the actor "Ego" (the acting agent) be sure to be understood by another actor "Alter" (the reacting agent)? A solution to this problem is at hand if we treat the symbols of a language as generalized forms of complementary expectations between social actors. From this standpoint to use symbolic expressions in interaction is to have certain expectations about how the addressee has to react if he understands. For example, when we use the word "red" we expect all members of the language community to identify the same objects as being red. This is what is meant by the characterization of language as a "stock of generalized expectations".

A common language is the main resource in social interaction if the problem of double contingency concerns the meaning of a situation. As long as a situation can be described in terms of a common language for each actor, the expectation that for all agents involved the situation has the same meaning would seem to be justified. Obviously, such a form of common understanding is a prerequisite of all social coordination. But to understand each others actions does not imply that Ego and Alter are motivated to interact in a coordinated manner. Thus, in addition to the common language human societies use, a number of specialized "languages", intended to enhance the probability of coordinated interaction, have been developed. For example the language of money that deals with problems of allocation or the language of power that deals with the problem of how to generate collectively binding decisions.

What is the difference between the common language and these generalized symbolic media of interaction in the narrower sense? Each symbolic utterance Ego addresses to Alter is an prompt to look at the world from a particular point of view: Thus Ego may offer an apple in order to obtain money in return; or Ego may allow Alter to take any apples except those from a certain tree; or Ego may point to her superior expertise and advise Alter to choose a particular apple by promising him God-like powers once he has tasted it; alternatively Ego may ask Alter to give her the apple as an act of compassion. The main question here is, why should Alter adopt Ego's selection? The answer leads us to the four generalized media of interaction Talcott Parsons has identified: money, power, influence and value commitment (Parsons 1969).

These media of interaction differ from the common language by promising special advantages or disadvantages that are bound up with adopting or rejecting the proposed selection. In the first of our four examples Ego is motivated to give Alter the apple by the expectation that this will give

her the possibility of buying something else. This expectation is mediated by the circulation of money. In the second example Alter's reaction will be influenced by the expectation that Ego is able to punish him. Here, power is the medium that structures interaction. In the third example the interaction is mediated through the expectation that Ego has superior knowledge or is competent for other reasons so that it is to Alter's advantage to adopt her world view. Parsons calls this medium "influence". And in the fourth example Alter's motivation to react as Ego suggests lies in the intrinsic value of this reaction. In this case, value commitment is the medium of interaction.

We don't want to claim that this is the complete list of generalized symbolic interaction media nor that it is the only possible way to classify them. Taking Parsons as a starting point, Niklas Luhmann, for example, has developed a concept of media of communication that places truth and love alongside money and power as the most elaborated media in modern societies (Luhmann 1975). Moreover, Luhmann further suggests there may be other media such as religious faith, art or basic civil values. What we are interested in here is how these or other media help to structure processes of coordination. The general answer is they perform in this way by enhancing the probability that the reacting actor will react as expected. In making use of one of the media of interaction Ego does two things at once: (1) She proposes looking at the situation from a particular point of view, a view that labels certain (re)actions as morally right or wrong or as allowed or forbidden by the powers that be, certain information as true or false or certain objects as interchangeable goods. (2) In framing or defining the situation in the way suggested, Ego offers Alter a particular motive to react according to these particular selections.

Thus, generalized media of interaction reduce uncertainty in two ways. Firstly, they simplify processes of coordination enormously. By using one of these media Ego and Alter are able to avoid a lot of coordination overhead that otherwise would make mutual adjustment much more complicated. For example, in an interaction mediated through money the coordination problem is reduced to a matter of goods and prices. Ego and Alter don't have to become friends if they want to exchange goods, they don't have to come to an agreement about religious beliefs, tastes or value commitments. The only problem they have to solve is to balance the price Ego is willing to pay with the price Alter wants to bring. Secondly, media of interaction reduce uncertainty by making each agent's actions predictable to a certain degree. Thus media allow Alter to generate expectations about what would follow his reaction; at the same time they allow Ego to generate expectations about the expectations Alter entertains. Just as Alter's expectations influence his reactions, so Ego can use his expectations about Alter's expectations in order to anticipate them.

The media of interaction are symbols of generalized expectations. If we use money to buy or sell goods, we do not interchange different goods but exchange goods with the generalized expectation of buying something else. Modern

money has no intrinsic value. It is nothing else than the symbolic form of such a generalized expectation. To accept money is to trust that each future transaction will come up to this expectation. Otherwise, there would be interchange merely between inherently valuable goods. The same applies to power, influence, and so on. Power symbolizes the expectation that an actor will be able to punish disobedience. Without this expectation power would be reduced to physical coercion and would lose most of its capacity to coordinate actions. Influence symbolizes the expectation that an actor possesses qualifications that make him an expert in the field in question. To adopt an actor's suggestions because of his influence is to trust in his superior skills. Otherwise, everybody in need of a certain expertise would have to acquire it on his own. And value commitment symbolizes the expectation that there are certain actions each actor will do or refrain from doing because they are morally right or wrong. Coordination based on altruistic behavior would be impossible without this expectation.

Like language in general, the special languages mediate the development and stabilization of complementary expectations because they themselves are generalized symbols of these expectations. And like the common language they also represent social conventions. That is, the expectations of Ego and Alter about what will follow next if they refer to one of the media and their motivations for reacting in a certain way are not merely a matter for individual choice. Rather, they keep to the beaten tracks, tracks that have been made by a long procession of previous interactions. To recognize how effectively social interaction is structured, you only have to imagine how to solve your daily coordination problems without the aid of these media.

Generalized Media of Interaction in Inter-Agent Coordination

In our opinion the concept of generalized symbolic interaction media is a very useful tool when it comes to tackling the phenomenon of social intelligence. It shows that many of the intelligent solutions for our daily coordination problems are not primarily the result of individual genius but stem from the fund of socially established expectations we routinely draw upon. Much of our social expertise consists in the ability to use the media of interaction our society provides in order to generate complementary expectations about each other and to act accordingly. So if we are interested in the question of how to feed artificial agents with social expertise and if we select the social expertise of human actors as our model, it would seem reasonable to use the media of interaction as a means of inter-agent coordination.

In DAI the problem of double contingency often is treated as a problem of how to reduce the amount of knowledge needed if each agent must take into account every possible combination of the whole range of agents' possible actions to achieve coordinated interaction (Durfee 1995). This way to address the problem of double contingency starts with the assumption that the situation in which

the coordination problem occurs is defined by a fixed number of agents where each of them is capable of a limited number of alternative actions. From this standpoint the problem of coordination is how to select from all possible combinations of actions the one combination that will be supported by all agents. Furthermore, given that each actor's choice is influenced by each of the other agents' choices, and even assuming a very limited number of agents and a highly restricted range of possible actions, agents still have to know a great deal about each other before they can successfully coordinate their actions. This becomes even more applicable when we leave "closed-world assumptions" behind and move out onto large-scale open systems (Hewitt 1991; Gasser 1991).

Providing agents with a common language and allowing them to communicate facilitates the acquisition of knowledge about each other, about each others knowledge, and so on. But a common language still doesn't help to reduce the possibly relevant knowledge to "knowing just enough to coordinate well" (Durfee 1995: 406). Instead, the problem of how to reduce knowledge now resurfaces as the problem of how to reduce communication. DAI has long since responded to this challenge with a number of practical ideas on how to reduce the amount of communication needed in inter-agent cooperation. Concepts like blackboard architectures, "action-at-a-distance" (Gasser 1991: 116) or "remote referencing" (Hewitt and Inman 1991) must be viewed in the line of practical containment strategies. Similarly, merging predefined cooperation primitives has been suggested to reduce communication overhead (Lux and Steiner 1995: 266). Other proposals include minimizing interdependence relations, conceptualizing coordination without communication and implementing patterns of interaction which operate on communication-free principles (Genesereth, Ginsberg, and Rosenschein 1986). Most of these suggestions are ad-hoc inventions with little rationale behind them beyond the intuitive demand for some technical device that prevents agents from talking too much.

From a sociological point of view, however, this underlying intuition is well worth considering because of an interesting parallel it presents with the way human societies handle the problem of double contingency: before ever becoming acute, the problem is solved by preexisting social structures. In human societies generalized symbolic media of interaction serve as a means to simplify coordination processes in a similar way to that used by mechanisms for technically reducing communication overhead in inter-agent coordination. But on closer examination there is a crucial difference between these two solutions: Generalized media are resources that only exist as long as they are continuously used to generate complementary expectations whereas the mechanisms used to reduce inter-agent communication are mostly subject to predefined restrictions. Basically, these mechanisms represent strategies to limit an agent's choices in a given situation, either by predefined global structures like social laws, common tasks and fixed structures of interchange or by predefined programs of

action for the individual agents such as fixed social roles, forced benevolence and so on.

Ensuring that coordination works this way runs the risk of reducing agents' actions to predefined patterns and of degrading coordinated interaction to the status of "pre-established harmonies" (Castelfranchi 1990: 50). This may be an adequate solution for coordination problems if the structured interaction concepts are derived from insect societies. But if we are interested in agency as a generating mechanism for true socially intelligent coordination, then we cannot be satisfied with agents which only react according to given individual or global structures. Seen in this light, the problem of coordination has two distinct facets: on the one hand, there is a need for mechanisms that facilitate coordination processes if we want to avoid renegotiating coordination solutions each time anew, reconsidering all possible solutions. On the other hand, however, if these mechanisms are designed as predefined structures that force the agents to act in a certain way, then we will end up with agents without agency. In our opinion, the concept of generalized symbolic media of interaction provides a means of escaping from the horns of this dilemma.

First and foremost from the point of view of generalized media of interaction, the time-honored controversy as to whether social agents should be friendly agents, who wish to do what they are asked to do or egoistic agents, only cooperating to improve their own self-interest becomes obsolete. In human interaction, benevolence or selfishness are not the fixed properties of individual actors but rather describe the behavior of actors in interactions mediated through value commitment or, for example, through money. In other words, whether one is a friendly or a selfish actor does not primarily depend on an individual predisposition. Rather, it depends on how a group of actors commonly perceives a situation with reference to one of these generalized media. Each action that refers to money, power, influence or value commitment communicates a suggestion as to how in that situation the actors should view themselves and each other. And it is the adoption of such a suggestion that makes them being helpful or selfish, obedient or influential. In this way, the generalized media serve as resources that facilitate coordination processes without imposing predefined limitations on the individual actors' programs of action.

Designing coordination mechanisms between artificial agents by trying to emulate the way money, power, influence or value commitment mediates human interaction is not a completely new idea. There are, for example, a number of contract net approaches that use a kind of money to coordinate competing agents (Wellman 1992). Trying to achieve a cooperative behavior in agents through social laws (Shoham and Tennenholtz 1992) is an approach that refers to power as the medium of interaction. The distributed vehicle monitoring testbed is an example of how the agents' particular expertise can structure coordination (Lesser and Corkhill 1983). In this case, influence can be viewed as the underlying medium. And, as we have seen, the benevolence assumption represents a sort of value

commitment. But once again, if we compare these concepts with our notion of generalized symbolic media of human interaction, the main problem is that they are imposing predefined limitations that reduce agency rather than providing resources that enhance the social expertise of artificial agents.

To arrive at a more dynamic view of structuring mechanisms in inter-agent coordination based on the concept of generalized media we have to examine the following relations more closely: (1) the relation between the generalized symbolic media as part of the preexisting culture into which the individual is born and as cultural patterns that only exist as long as they are continuously reproduced by the individuals' actions. (2) The relation between the structuring capabilities of generalized media and the agents' capability to act differently. (3) The relation between the various generalized media as alternative resources to facilitate coordination processes.

The common language and the more specialized languages of money, power, and so on are social conventions that predate each individual. This may lead us to believe that they are external structures forcing individual actors to act in a specific way. But this is only a half-truth. It is true that the generalized media structure sequences of interaction by providing stable expectations about the behavior of the actors in particular situations. But they do so in a kind of a self-fulfilling prophecy: such expectations are justified only as long as they are commonly maintained by the individual actors. That is to say, the beaten tracks of the generalized media remain beaten tracks only as long as they are continuously followed. Attempts to model coordination mechanisms between artificial agents as predefined structures ignore this interdependency between agency and structure. We should rather be looking for ways to design coordination mechanisms that are both a precondition and a result of the interactions they mediate. Coordination by generalized expectations in our opinion is a step in this direction: Generalized expectations are always part of the preexisting patterns of an agent society that agents can draw upon to coordinate their actions. But since no agent can be completely sure whether the other agents' actions will meet his expectations, and since this will influence his future use of the generalized expectations they are at the same time a result of the interactions they mediate.

This leads on to our second point. In human interaction there is no guarantee that the actors will behave as expected. Generalized media communicate expectations about the advantages or disadvantages concomitant on certain actions. But ultimately it is up to the agents themselves to adopt or to reject the choices suggested in this way. Unless a certain action is directly forced by physical coercion human actors are always able to refuse to obey other actor's orders, to reject proffered goods, to disappoint friends, and so on. Accordingly, if the generalized media are to be our model of coordination mechanisms between artificial agents, then we have to be careful not to construct structures of interaction that leave agents no choice but to act cooperatively. We should rather view these coordina-

tion mechanisms as resources to generate complementary expectations between agents, aiming to motivate them to choose from the range of actions open to them the very one they are expected to choose. From this point of view, it must also be allowed that these expectations may fail to motivate an agent, that he may have other preferences and that he may act in a different way.

As long as we are concerned with artificial agents whose actions are coordinated by only one of these generalized media, the distinction between rules that force and expectations that motivate agents to act in a certain way seems to be academic indeed. With only one way of generating expectations as to the possible advantages or disadvantages attendant on certain actions, there is no reason why an agent should be motivated to act unexpectedly. But this scenario changes completely if we introduce a variety of ways of generating expectations from which different preferences to act can be derived. As human actors in modern societies, an important part of our social expertise is based on the capability to distinguish between the different forms of generalized expectations the society provides, to initiate and to understand social interaction by selecting between such generalized expectations and to evaluate the advantages and disadvantages of our possible actions in situations that are mediated through one or another of them. It is this expertise and not an inability to act differently that allows us to interact in a manner that is commonly perceived to be adequate. Thus, in artificial societies the development and use of different forms of generalized expectations is a vital prerequisite for installing a similar kind of social expertise in artificial agents.

From Social Agents to Artificial Societies

In our paper we have given a rough idea of how to tackle the issues of agency and society from a sociological perspective. In taking social interaction and the problem of double contingency as the starting point of our analysis, we have outlined how generalized media of interaction contribute to large-scale coordination in modern societies by shaping and structuring the very process of social interaction from which they emerge. Generalized media of interaction must be considered as one of the most powerful mechanisms in modern human societies for achieving coordinated action. In co-evolution with other mechanisms (differentiation, institutionalization etc.), they maintain coordination at a level of complexity far in excess of the restricted capabilities of insect colonies, ape communities or prehistoric human tribes. What we suggest is framing the concept of interaction media as a source of inspiration from which to design and to build artificial social systems with features that take account of the enormous complexity, robustness, and capacity for innovation of human societies. Admittedly, human societies are a particularly hard case for computer scientists to crack. Subhuman societies or communities are much easier to model and to implement because of their relative simplicity; complexity always poses a compelling argument for refraining from implementation.

But the challenge still remains: can we develop computational concepts in DAI or in other fields of research related to socially intelligent agents capable of coping with the hardest possible case - with modern society?

Of course when we ask questions about how to design artificial societies and how to make sociological categories computational without loss of complexity, we immediately raise a whole host of other questions. It is not the purpose of this paper to answer these questions but rather to draw attention to the sociological difference between (micro) "community" and (macro) "society", for it is along such lines that questions must be posed and, eventually, answers must be given. Since for the most part they are innocent of the distinction between society and community, researchers in DAI and other fields of artificial social agents "naturally" seem to prefer micro-approaches. We can, however, observe a growing awareness of the problematic aspects of micro-approaches to multi-agent systems, as for instance when DAI researchers are confronted with an exploding communication overhead. In fact, modern human societies would suffer instantaneous breakdown if they had to maintain coordination solely on the principles of direct communication, personal acquaintance and physical co-presence alone. By way of contrast, coordination in ant colonies, wolf packs, gorilla groups is based on kinship and must be maintained by permanent face-to-face-interaction and direct contact (e.g. grooming). Physical co-presence is needed both to confirm and to (re)establish affection and solidarity among group members, as well as to define and redefine an individual's hierarchical position and social role vis-à-vis the other members of the community. Similarly, social coordination in prehistoric human tribes is based on kinship relations, but kinship here is gradually transformed from a natural into a social relationship. Indeed, kinship is eventually emerging as a genealogical value system that can be construed as one of the first ancestors of generalized media of interaction.

Coming back to modern societies and DAI: how is "scaling-up" possible from small paleolithic tribes of hunters and gatherers to differentiated superstructures of the highest complexity? How can we attain global coherence at the macro-level without serious obstruction from the concurrent actions of millions of autonomous individuals following their own goals and prerogatives? If we do not want to retreat to the micro-level of small communities of reactive, embodied, or knowledge-based agents, we must try to specify the "link" between micro and macro; we must show how macro-structures and micro-structures are functionally integrated, and how generalized media of social exchange like value commitment, money, power, and influence, are merged with, and mediated by, acquaintances, personal relationships, face-to-face-interaction, social grooming etc. If DAI wants to construct large-scale artificial societies, these questions should be placed at the very top of the agenda.

References

- Castelfranchi, C. 1990. Social Power. A Point Missed in Multi-Agent, DAI and HCI. In *Dezentralized AI, Proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 49-63. Amsterdam u.a.: Elsevier Science Publishers.
- Drogoul, A., and Ferber, J. 1994. Multi-Agent Simulation as a Tool for Modeling Societies: Application to Social Differentiation in Ant Colonies. In *Artificial Social Systems. 4th European Workshop on Modelling Autonomous Agents in a Multi-Agent World. MAAMAW '92*, 3-23. Berlin u.a.: Springer-Verlag.
- Durfee, E. H. 1995. Blissful Ignorance: Knowing Just Enough to Coordinate Well. In *ICMAS-95. Proceedings of the First International Conference on Multi-Agent Systems*, 406-413. Menlo Park u. a.: AAAI Press/ The MIT Press.
- Gasser, L. 1991. Social Conceptions of Knowledge and Action: DAI Foundations and Open Systems Semantics. *Artificial Intelligence* 47: 107-138.
- Genesereth, M. R., Ginsberg, M. L., and Rosenschein, J. R. 1988. Cooperation without Communication. In Bond, A. H. and Gasser, L. eds., *Readings in Distributed Artificial Intelligence*, 220-226. San Mateo, Ca.: Morgan Kaufmann Publishers.
- Giddens, A. 1984. *The Constitution of Society: Outline of the Theory of Structuration*, Cambridge: Polity Press.
- Hewitt, C. 1991. Open Information Systems Semantics for Distributed Artificial Intelligence. *Artificial Intelligence* 47: 79-106.
- Hewitt, C., and Inman, J. 1991. DAI betwixt and between. From "intelligent agents" to open system science. *IEEE Transactions On Systems, Man, and Cybernetics* 21(6): 1409-1419
- Lesser, V. R., and Corkhill, D. D. 1983. The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem Solving Networks. *The AI Magazine* 4: 15-33.
- Luhmann, N. 1975. Einführende Bemerkungen zu einer Theorie symbolisch generalisierter Kommunikationsmedien. In Luhmann, N., *Soziologische Aufklärung 2. Aufsätze zur Theorie der Gesellschaft*, 170-192. Opladen: Westdt. Verlag.
- Lux, A., and Steiner, D. 1995. Understanding Cooperation: An Agent's Perspective. In *ICMAS-95. Proceedings of the First International Conference on Multi-Agent Systems*, 261-268. Menlo Park u. a.: AAAI Press/ The MIT Press
- Parsons, T. 1969. *Politics and Social Structure*, New York: Free Press.
- Parsons, T., and Shils, E. eds. 1951. *Toward a General Theory of Action*, Cambridge, Mass.: Harvard University Press.
- Shoham, Y., and Tennenholtz, M. 1992. On the Synthesis of Useful Social Laws for Artificial Agents Societies (Preliminary Report). In *AAAI-92. Proceedings Tenth National Conference on Artificial Intelligence*, 276-281. Menlo Park: AAAI Press/The MIT Press.
- Wellman, M. P. 1992. A General-Equilibrium Approach to Distributed Transportation Planning. In *AAAI-92*, 282-289. Menlo Park u.a.: AAAI Press/The MIT Press.