JASSS

Thomas Malsch and Ingo Schulz-Schaeffer (2007)

Socionics: Sociological Concepts for Social Systems of Artificial (and Human) Agents

Journal of Artificial Societies and Social Simulation vol. 10, no. 1
<http://jasss.soc.surrey.ac.uk/10/1/11.html>

For information about citing this article, click here

Received: 21-Dec-2006 Accepted: 21-Dec-2006 Published: 31-Jan-2007

Abstract

Socionics is an interdisciplinary approach with the objective to use sociological knowledge about the structures, mechanisms and processes of social interaction and social communication as a source of inspiration for the development of multi-agent systems, both for the purposes of engineering applications and of social theory construction and social simulation. The approach has been spelled out from 1998 on within the Socionics priority program funded by the German National research foundation. This special issue of the *JASSS* presents research results from five interdisciplinary projects of the Socionics program. The introduction gives an overview over the basic ideas of the Socionics approach and summarizes the work of these projects.

Keywords:
Socionics, Sociology, Multi-Agent Systems, Artificial Social Systems, Hybrid Systems, Social Simulation

---

What is Socionics?

1.1

Socionics is a combination of sociology and computer science. As a term Socionics was coined in 1996 (Malsch, Florian, Jonas & Schulz-Schaeffer 1996, Malsch 1997) to attract public and scientific attention to a promising new research field outside the confines of what Kuhn used to call the "normal sciences". The general idea of Socionics can be compared to that of Bionics. Bionics is the attempt to model new technologies on biological phenomena. As an example take the Velcro or hook-and-loop fastener modelled on the fruit of the cocklebur (latin: *arctium lappa*), whose original function is biological dissemination by getting parasitically stuck to the fur of a by-passing animal. Socionics comes forward with a similar proposal: to exploit sociological conceptions as a source of inspiration for the development of distributed software such as multi-agent systems (MAS) or artificial social systems (ASS), both for the purposes of engineering applications and of social theory construction and social simulation.

1.2

Research in Socionics began around 1980 when American computer scientists coming from the field of Artificial Intelligence (AI) turned to distributed problem solving and collaborative intelligence and became involved with sociologists who on their part wanted to understand and explain the social and organizational impacts of computation. At that time sociological studies on artificial social systems and social simulation deploying techniques from Distributed Artificial Intelligence (DAI) on one hand, and DAI research importing sociological conceptions into agent-oriented software engineering on the other, was still in its infancy.[1] In the early 1990s the new research field, which so far had hardly been noticed by mainstream AI and sociology, began to make progress, and with the awareness growing that DAI had a "problem of society construction" (Müller 1993), Gasser's dictum was eventually acknowledged that "current AI is largely a-social, and because of this, it has been inadequate in dealing with much human behaviour and many aspects of intelligence" (Gasser 1991: 108). As a consequence, at the end of the 1990s a wave of research activities was initiated.

1.3

In 1998 the German National Research Foundation (DFG) inaugurated the Socionics priority program (Schwerpunktprogramm Sozionik, SPP 1077), with funding provided for six years. Initiated by a group of researchers from both disciplines,[2] the Socionics priority program's intention was to promote and encourage interdisciplinary research collaboration between sociologists and computer scientists (Müller, Malsch & Schulz-Schaeffer 1998; Malsch 2001). In contrast to most previous work in DAI, the point of departure for Socionics was social theory and sociological observation. It is from sociological knowledge and not from everyday social metaphors or naive images of sociality that the Socionics research projects seek to conceptualize the distributed behaviour of multi-agent systems. In order to ensure verifiable interdisciplinary cooperation, an appropriate institutional structure was given to the Socionics program: all applicants who wished to join were obliged to form their own bi-disciplinary tandem projects with shared common responsibility. This was a condition sine qua non to apply for funding. Despite the inbuilt hurdle of shared tandem responsibility, more than twenty tandems comprising over forty research teams responded when the Socionics call was published. Out of these applications, six tandem projects of sociologists and computer scientists with twelve research teams from different universities (Aachen, Bamberg, Berlin, Chemnitz, Hamburg, Konstanz, Munich, Saarbrücken) were chosen to join the Socionics priority program.

Socionics and the problem of social order

2.1

When the first projects were launched in October 1999 no one was actually able to predict the boom which Socionics and related work on artificial agent societies would experience in the coming years. Today research and design of artificial social systems is widely recognized as one of the fields of major future engineering interest. In an influential paper of the Feldafinger Kreis, a joint initiative of the German Industrial Federation (BDI) and the Fraunhofer-Gesellschaft (FhG) of Engineering Research, Socionics is explicitly proposed for public funding and institutional support (Wahlster & Weyrich 2002: 15). Meanwhile the European Community has initiated related R&D programmes such as "AgentCities", "AgentLink II", "Future and Emerging Technologies" (FET)[3] and has explicitly suggested to promote knowledge transfer "between agent technologies and the social sciences" (Luck et al. 2003: 25ff). Also a diversity of scientific conferences and interdisciplinary journals such as the *JASSS* or *Autonomous Agents and Multi-Agent Systems* were founded and a growing number of publications on Agent Societies, doctoral theses and monographs (Ossowski 1999, Verhagen 2000, Hales 2001), special editions of scientific journals (Conte et al. 1999, Edmonds & Dautenhahn 2001) and conference proceedings (Ahrweiler & Gilbert 1998, Prietula, Carley & Gasser 1998, Omicini, Tolksdorf & Zambonelli 2000) have appeared. Today we can say that the formerly exotic idea of importing sociological concepts into DAI seems to be widely accepted (Panzarasa & Jennings 2001, Zambonelli & Parunak 2003).

2.2

With hindsight it is interesting to note that the Socionics program was indeed lucky to catch the wave peak at the right moment and to benefit from surfing on an international scientific development which was strongly propelled both by the rise of the New Economy and new challenging research questions. However, this does not really explain why artificial social systems became an interesting topic and what kind of common issues were actually at stake that motivated sociology and DAI to cooperate. In this context it may be instructive to recall the motives and goals of classical AI. AI regards the human mind or the human brain as the locus of intelligent problem-solving. Hence, AI's research seeks to construct technological equivalents of the cognitive skills of the rational single human being. In contrast to AI, however, DAI assumes that intelligence is rather a social phenomenon and that complex problem solving is inherently based on decentralised action and distributed cooperation of a multitude of autonomous intelligent units. DAI wants to replicate the coordinated social intelligence of a multitude of artificial agents for engineering purposes. And this means that modelling social systems is not regarded as a goal in itself (as it is for sociology and social simulation) but a means to an end.

2.3

While sociology and social simulation are interested in multi-agent systems because these "seem more able to 'mirror' societies and groups of people than their alternatives" (Doran & Gilbert 1994: 10), DAI is not primarily concerned with trying to understand what makes human societies tick but rather to extract from the study of social phenomena computable principles of decentralised control architectures. Hewitt's critique of the "closed world assumption" of classical AI nicely illustrates this point: If it is true that complex problems usually cannot be represented as closed worlds by just one single "microtheory", but need to be expressed by at least two "microtheories", the consequence is "logical indeterminacy" (Hewitt 1977). In logical terms this problem is unsolvable. As it turns out, the only way out of the dilemma of Hewitt's logical indecision problem is negotiating conflicting microtheories at a meta-level of abstraction. Thus "contract" and "negotiation" (Davis & Smith 1983) were the first social metaphors to be adopted by DAI as an early attempt to manage dissent and conflict in a socially intelligent way. Since such principles cannot be invented merely by common sense or "computational introspection" alone, it is at this juncture that DAI was looking for new impulses from the organizational and social sciences (Bond & Gasser 1988, Huhns & Singh 1998).

2.4

That DAI, in granting individual autonomy to intelligent agent programs, runs into massive control and coordination problems, is not really astonishing to sociologists, since these problems can be compared to the classical question of social order: How is it possible to maintain social order in the face of intelligent agents that are capable of autonomously taking self-interested decisions? In DAI terms the problem of social order has been stated by Jennings as follows: "Agents may spread misleading and distracting information, multiple agents may compete for unshareable resources simultaneously, agents may unwittingly undo the results of each others activities and the same actions may be carried out redundantly... the dynamics of such systems can become extremely complex, giving rise to nonlinear oscillations and chaos. In such cases the coordination process becomes correspondingly more difficult as well as more important" (Jennings 1993: 227). Indeed, the "price of anarchy" (Papadimitriou 2001) which must be paid for the construction of open artificial societies is high compared to that of closed multi-agent systems. Openness is difficult to manage both technically and socially because this implies allowing open access for uncontrollable agents.

2.5

This problem is particularly challenging for those who want to build "open multi-agent societies" (Dellarocas & Klein 2000, Artikis & Pitt 2001, Davidsson 2001). Consider, for instance, the case of electronic market places. These are open to and visited by huge numbers of unknown mobile agents and so designers, operators and users are faced with a complex bundle of safety and security risks such as malevolent or unintentional overstraining of computational resources or destruction of data and code. In order to keep these risks under control, a variety of proposals has been discussed that reach from "partially controlled multi-agent systems" (Brafman & Tennenholtz 1996) to "semi-open" or "semi-closed systems" (Davidsson 2001): (i) no access controls but in-built compliance with normative rules of benevolent behaviour (Wellman & Wurman 1998, Gustavsson & Fredriksson 2001); (ii) no normative rules of behaviour but access controls by a gate keeper agent (Davidsson 2001); (iii) neither normative rules nor access controls but real-life authentification and traceability of each agent back to its owner, that is to a natural or legal person who carries full responsible for his or her agent's actions (Davidsson 2000, Verhagen 2000).

2.6

Thus, unwanted dysfunctions could be coped with by introducing strict regulations combined with negative sanctions and preventive security checks. However, this is not as easy as it looks. There are two main reasons why control strategies implementing restrictive social rules are deeply problematic for MAS engineers. The first is that excessive controls are poison for applications such as electronic marketplaces which cannot function without a certain amount of unpredictable individual action. The second is that even if all agents were benevolent and compliant with social norms, it is impossible to completely out-design fault, failure and disruption. As Dellarocas and Klein have argued, the tacit engineering assumption of most system designers in DAI, according to which "the emerging society will exhibit stable and efficient behaviour... if all agents follow the 'right' mechanism", is essentially doubtful (Dellarocas & Klein 1999: 2). Even if all agents are strictly subjected to screening procedures or other access control techniques and do indeed interact faultlessly with the best of cooperative intentions, unintended dysfunctions may arise at any time at the system task level of an artificial agent society.

2.7

As soon as MAS are scaled-up with larger and larger amounts of interacting agents, non-intentional effects such as resource poaching or distractions from tasks and responsibilities cannot be avoided (Chia et al. 1998). The basic reason is that systems dynamics become more and more unmanageable with growing interactive and temporal complexity. In order to understand the dynamics of complex systems it may be useful to draw on concepts such as "self organisation" or "emergence". Swarm formation, a collective behaviour observed in large collections of birds or fish is a canonical example for non-equilibrium phenomena at the global or system level (Bonabeau et al. 1999). At the systems level swarm behaviour exhibits an emergent new quality of highly coordinated collectively movements, which is essentially different from the individual performance at the component level, despite the fact that those patterns of collective behaviour at the system level are made up of nothing other but of the behaviour of its individual components.

2.8

What makes theories of emergent self organization so difficult to apply to complex social interaction is the fact that the basic "components" of a social system are knowledgeable, sense-making individuals. This does not only mean that these individuals are inherently complex by themselves but that they are capable of reflexively anticipating and even outwitting the outcome of collective social interaction at the global level of social structure formation. It is beyond question that social theories of emergence and self organisation must be regarded as highly attractive design principles for an agent-oriented software engineering (Kephard & Chess 2003) interested in constructing agent societies "that self-organize... in response to task requirements" (Klusch 2001). But how can this goal be achieved without taking pre-existing social norms and regulations into account?

2.9

The difficulty is that these principles cannot be easily translated into programming techniques of collective activity that emergently can produce a set of negotiated social contracts or social norms on one hand, when we have to acknowledge on the other that these norms are, in turn, enforced through mechanisms of social institutions which must be there already before any contracts can be negotiated. After all it is not really surprising when Mamei and Zambonelli must admit that "we (the DAI community) still do not know how to program and manage... autonomous self-organizing systems. The main conceptual difficulty is that we have direct engineering control only on agent's local activities, while the application task is often expressed at the global scale." (Mamei & Zambonelli 2004)

How to model social order

3.1

What we can learn from our tour de force through MAS and AAS research is that socially-based software engineering is confronted with questions quite similar to those which sociologists seek to answer. As we can see from the following presentations of Socionics tandem projects, both disciplines have a common saddle point which is called the micro-macro-problem or the old question of social order. This can be translated into the question of how to model the relationship between social (inter)action and social structure (social roles, regulations, institutions). It directly refers to the sociological controversy whether social structure should be explained "bottom up" as an emergent outcome of social action or whether social action should be conceived of as being constituted "top down" from social structure. In DAI research the sociological controversy appears under the guise of designing types of coordination or social laws either "offline" by the systems engineer or "online" by the agents themselves at run-time. The papers collected in this special section seem to demonstrate that neither "bottom up" nor "top down" approaches can be realised in a pure form. Rather, the papers can be read as searching for different ways to combine and integrate both approaches, allowing for different degrees of freedom of social exchange bound to different degrees of normative prescription. This has to do with but is not actually determined by the different sociological appraoches chosen by the tandem projects.

3.2

However, modelling becomes more complicated if on top of the two dimensions of (i) agent interaction at the local level and (ii) social regulations at the global level a third dimension is added, namely (iii) engineering task accomplishment at the global level. Task accomplishment is specifically interesting for computer scientists, while sociologists primarily want to explain the relationship between dimensions (i) and (ii). Again, task accomplishment at the global level can be modelled as the emergent outcome of autonomous agents interacting at the local level or as a largely predetermined result of restrictive regulations and rules of action. In both cases, in simulating how social structures are shaped as well as in simulating how global tasks are performed, the structure of explanation seems to be quite similar. What is more, the models and prototypes presented here obviously can be tested and validated both as sociological tools and as engineering technologies.

Modelling dynamic message reference networks

4.1

There are two different ways of analysing the structural features of a social system. One way is to look for behavioural regularities of the individual entities that make up the social system and for regularities with respect to the ways the entities interrelate their individual behaviours. This is the way we are most familiar with. The other more unfamiliar way is to scan the traces the social individuals leave behind for regularities. The COM/TE approach (Malsch et al., in this special section) follows the second strategy. The traces left behind which the researchers are focusing on are empirical message signs in whatever form — books, articles, emails, chat postings, utterances etc. The social structures they are analysing are the relational patterns between massages that dynamically emerge from references between messages. Niklas Luhmann's view of social processes as chains of communication events and of social structures as patterns of communication of the same kind provides the conceptual background of the project.

4.2

In particular, the COM/TE project looks at patterns of reputation. Reputation is operationalised as the social visibility of a message sign. Social visibility is designed as a variable that measures the number of references a certain message gets from other messages. This reflects the assumption that messages become visible through other messages referring to them. The variable is also influenced by the number of reference the message itself sets to other messages. The idea here is that a reference to a visible message enhances the visibility of the message setting this reference, too. Additionally, social visibility includes an aging factor which is intended to model the fact that messages by growing older may get lost or may become less interesting. With operationalising reputation in this way the COM/TE researchers adopt from bibliometry and scientometry the assumption that counting citations which occur within a certain time period is a suitable way to grasp the influence of a paper or of a message in general. On the other hand, they take into account that qualitative features are also important, as for example the topic and the relevance of a message or its author. However, features of this kind are not implemented in the basic version of the COM/TE approach presented below.

4.3

For the means of simulation message reference networks the social visibility variable is presented as a dynamic function. By manipulating the parameters of this equation (i.e. the default visibility of each message, the number of in-going and out-going references, and the speed of aging) the COM/TE researchers are able to produce three distinct patterns of message referencing. They call these three patterns the modernist style, the classicist style, and the historicist style of communication. The historicist style is characterised by the fact that some messages which have gained a certain visibility within the first few generations of the communication process become and remain to be the most visible messages throughout the whole process. The historicist style, thus, represents a pattern of communication that heavily relies on foundational messages. Theological communication that refers to the bible or historical publications citing ancient sources are the authors' examples for this style. The classicist style shows a pattern where the most visible messages are more or less evenly distributed over the whole sequence of the communication process. Within every generation of the communication process messages do have an equal chance of becoming visible. According to the authors, publications in sociology is an example of this communication style in that each generation of sociologists produces books or papers that turn out to become classical works through other researchers' citations. In contrast to the historicist and classicist patterns the modernist style enhances the visibility of those messages which have been produced in the present or the near past. This communication style pays attention only to new messages, a pattern the authors view to be characteristic for citation practices in natural sciences and computer sciences.

4.4

The next step of the COM/TE project is to explore by simulation how these three ideal typical communication styles are influenced by different media of the communication process. The core idea of this research step is that two variables are sufficient to represent key features of different media forms: message persistence and reference/reception ratio. The idea is simple but striking: Some forms of communication such as face-to-face communication or mass media communication are based on messages with a very limited life span. One has to respond immediately. Later, the message is forgotten or out of date (low message persistence). For other forms of communication, however, such as science communication and archive-based communication in general, storage of messages and reference to stored messages is essential (high message persistence). With some forms of communication such as face-to-face-communication or science communication to receive (i.e. to listen to or to read) a message means to respond. In contrast, there are other forms of communication such as mass media communication where typically only small portion of received messages evoke responses and references to it (low vs. high reference/reception ratio). Thus, based upon the two variables a fourfold table can be constructed which can be used to categorise the different media forms of communication.

4.5

The simulation runs of the COM/TE project of the three communication styles under the four different conditions as represented by the fourfold scheme show effects that mainly go back to differences in the message persistence: the modernist style remains unaffected, the classicist style gets a modernist touch, and the historicist style under the condition of low message persistence is transformed into a modernist style but otherwise becomes even more historicist. The effects of differences in the reference/reception ratio are less visible in the simulations. All in all, the COM/TE project shows that many insights can be gained by looking at structural features of communication processes already without taking the actors of these processes — those who read or listen to messages and respond to messages with new messages — into account. To model these processes of reception and inception (responding) explicitly by developing a multi-agent system will be subject to further research of the COM/TE project.

Modelling the emergent semantics of communication

5.1

People orient themselves in the social world as well as in the natural world by attaching meaning to events, actions, and situations. As their most important means of interaction humans use symbols which represent as signifiers the meaning of the signified. The meaningful interpretations of the social and natural world and their symbolic representations constitute a human society's common stock of knowledge. This stock of knowledge can be relatively stable if the respective society (or societal subsystem) does not expose itself to yet uninterpreted events.

Traditional societies are of this kind or social settings where a central authority has the power to define and assert a certain system of meaning. However, if a society or a social system does expose itself to yet uninterpreted phenomena, when it is open for new members from the outside or when it allows its members to develop new behavioural patterns, the common stock of knowledge is subject to continuous changes and the constitution of meaning is an ongoing social process. Much work in multi-agent research presupposes the rather stable meaning structure of traditional or authoritative social systems, for example, by implementing social rules and normative rules the agents cannot but obey. With respect to open multi-agent systems, however, a different approach is necessary: an approach which allows meaning to evolve during the ongoing processes of interaction between agents. The aim of the Empirical Semantics project is to develop such an approach.

5.2

The Empirical Semantics approach (Nickles et al., in this special section) is based on a pragmatic or "consequentialist" concept of meaning stating that the expected consequences of a communication define the meaning of this communication. The consequences of communications are observable events that follow from communications, i.e. the reactions of the communications' addressees. Observed consequences of past communication processes by extrapolation and anticipation can be transformed into expectations of future consequences of present or future communications. For this reason, their meaning lies in their expected consequences. The Empirical Semantics approach implements this concept of meaning by modelling utterances as projections, that is, as assertions about expected future states.

5.3

From the point of view of the Empirical Semantics approach, thus, analysing and modelling the emergence of meaning means to analyse and to model expectation formation. This is done in two different ways, on the one hand from the sociological macro-perspective of Luhmannian social systems theory, on the other hand from the sociological micro-perspective of symbolic interactionism (George Herbert Mead) and framing theory (Erving Goffman). Within the Empirical Semantics approach the main difference between both perspectives is that in the first case it is the social system itself which produces and actualises an expectation structure that reflects the ongoing process of communications and following consequences. The social system is modelled as a macro-social observer of all communications and consequential events who provides the expectations drawn from these observations as an overall social structure. In the second case, the expectation formation is done by individual agents which are modelled to be intentional, deliberative, goal-oriented systems. In contrast to the macro-social observer, these individual agents' expectations are derived from the limited empirical knowledge accessible to them and the expectation building is governed by the bounded rationality of assessing only interaction processes that are relevant with respect to their goals.

5.4

From the empirical observation of all occurrences of communications and subsequent events the macro-social observer computes for each occurring combination the probability of the event to follow the communication. As a result, it provides a data structure that is called expectation network in the Empirical Semantics approach. The expectation network is a graph with the events (communication acts and physical actions) as the nodes and the transition probability as the edges between them. During the ongoing interaction process the macro-social observer actualises the expectation network by modifying the graph according to the observed empirical reality.

5.5

The individual agents as micro-social observers store their experiences in form of frames. The micro-social approach within the Empirical Semantics approach models frames as patterns of interaction episodes. The individual agent adds each new interaction episode it experiences to his repository either as an additional case of an already existing frame or as a new frame if the new experience cannot be subsumed under an existing frame. Additionally, the agents employ generalisation methods for generating more abstract frames from certain existing frames if possible. Each frame describes a trajectory of events whose structure is similar to a path within the expectation network. The whole set of frames represents the individual agents expectation structure. The agents select the frame that is most appropriate to the situation at hand and act according to the trajectory of events it describes. In this way the individual agents' expectation structure governs their behaviour.

5.6

The most far-reaching aim of the Empirical Semantics project is to come to an integrated approach that combines the micro- and the macro-perspective. Since the expectation concept is central to both perspectives the authors believe this to be an idea with a reasonable chance of success. To make progress in this direction they suggest to take the interaction trajectories generated by the individual agents' frame-building and frame-adapting processes as the empirical data from which the expectation network as the overall expectation structure of the social system could be derived. This is an interesting idea but it lacks to take into account the reverse direction of the social influences on individual sense-making. Adding this aspect could lead to the consequence of treating frames not only as individual meaning structures but as intermediate structures where individual experience and commonly shared knowledge merge.

Modelling reputation as a multi-level concept for open systems

6.1

Coordination of tasks is difficult when the participants possess little knowledge about each others' competence and willingness to comply with the work agreements necessary to get the task done. Work contracts and the contract law act as means to reduce uncertainties of this kind. However, coordinating work in this way requires designed sanctioning mechanisms which may be costly if one has heavily to rely on them. Moreover, it presupposes that the work of the agent within a principle-agent relation is controllable by the principle, what is often not the case. Thus, it would be advantageous and sociologically interesting to have mechanisms producing reliable expectations about other participants' performance sort of as a by-product of the ongoing interaction itself. Reputation is such a mechanism. With open multi-agent systems one faces similar problems of controlling agents by means of pre-designed sanctioning mechanisms. Open multi-agent systems are characterized on the one hand by the fact that agents are free to join or to leave the system and on the other hand by the fact that they may change their behaviour during runtime. For both reasons pre-designed institutions may turn out to be too static to deal with changes within the system. Thinking about emergent mechanisms of social order such as the reputation mechanism, thus, has become a growing topic of multi-agent research.

6.2

The Social Reputation approach (Hahn et al., in this special section) belongs to this strand of research. It starts with the notion of reputation as symbolic capital. In contrast to other resources, this kind of resource does not designate something an actor possesses directly such as money, skills, or knowledge. Rather, reputation sums up from what others think of an actor, how other actors evaluate an actor's qualities. Reputation becomes a resource when such evaluations are used to decide with whom to cooperate and whom to exclude. By introducing reputation as a sociological multi-level concept the Social Reputation approach aims at developing an elaborated model of reputation as a mechanism of flexible self-regulation. In many sociological concepts it has proved useful to distinguish between three levels of social aggregation: the micro-level of social interaction, the meso-level of social groups and organisations, and the macro-level of society as a whole. Corresponding with these levels, the authors distinguish between three levels of reputation: (1) micro-level reputation, called "image": the aggregated evaluations an agent derives from his own experiences with the agent in question and from what he gets to know about a few other agents' experiences with the target agent; (2) meso-level reputation, called "social esteem": evaluations of an agent which are derived from the reputation of a group or organisation the agent in question is member of; and (3) macro-level reputation, called "prestige": evaluations provided by an impartial institutional agency able to aggregate over time a large amount of the individual experience-based evaluations.

6.3

The performance of these three forms of reputation is explored with a multi-agent simulation where customer agents assign tasks to provider agents based on the providers' reputation values. The reputation values are built up and accommodated by constant updates of the experience-based evaluations during successive runs of task delegation processes. Reputation is measured with respect to two subjects: the trustworthiness of providers, expressing the providers' reputation of fulfilling contracts, and the credibility of informants, i.e. their reputation of evaluating agents according to their experiences. Since the aim of the Social Reputation approach is to explore reputation as a dynamic self-adapting mechanism of social order in open systems, the performance of the three versions of reputation is analysed with respect to destabilising events typical for open systems: opportunistic behaviour, changing participants, and emergent group formation. By and large, the simulation results indicate that prestige performs somewhat better than image in dealing with opportunistic behaviour (fraud of providers and lies of informants) and with changing participants (when new providers enter the setting). This is what should be expected, since prestige represents a more comprehensive set of experience-based evaluations compared to image. That the differences are not very large, however, is due to the fact that the simulation scenario's population is relatively small (90 providers, 30 customers). It corresponds with everyday experience that with such a small setting a separate agency is not really necessary for distributing the participants' evaluations about each other. The only major difference occurs with respect to group formation. In the simulation scenario providers are allowed to form groups for sub-contracting tasks assigned to them. When it comes to the question to which kind of reputation providers refer for identifying trustworthy providers as sub-contractors it turns out that prestige performs much better than the other forms of reputation. The reason is that the required level of trustworthiness for this decision is set higher in the simulation scenario than in the case of customer decisions. The amount of agents involved and the degree of the quality the reputation value is able to represent thus turn out to be important parameters influencing the respective performance of the different aggregation levels of reputation.

Modelling the self-organisation of social units

7.1

The ASKO project (Köhler et al., in this special section) provides an approach for modelling social processes of structuring and structuration. The approach relies sociologically on self-organisation theory and is realised with advanced Petri net technology. Petri net technology does not only serve as a means of translating insights from social theory into features of multi-agent systems. Additionally, the mathematical character of Petri nets is used to evaluate and control the consistency and integrity of the underlying sociological concepts. This is an experience common to all projects of the Socionics research program: The need for precise and consistent definitions and ontologies as a precondition of modelling and simulating social reality forces social scientists to formalise their empirical observations and conceptual assumptions accordingly. On the one hand, this imposes restrictions on social simulation since social reality includes contradiction and incoherence. On the other hand, however, the need to formalise sociological concepts provides an opportunity for enhancing their quality. Petri nets possess the advantage of high visibility and traceability of the processes modelled with this formalism. Thus, the formalism is especially suited for detecting hidden implications of the social processes for identifying inconsistencies and hidden gaps of sociological theories.

7.2

The ASKO project's main conceptual subject is the development of a general theory of social self-organisation based on similarities between existing sociological theories. According to this conceptualization, the constitutional dynamics of social self-organisation consists of three processes: social units relate themselves to their environment (practical acceptance), they ascribe to each other stable regularities of behaviour (structural capability), and they symbolise structures and capabilities to stabilise and visualise them (reflexive symbolism). The social units are assumed to be the agents as well as the results of this threefold process. The ASKO researchers argue that not only human actors but social structures and symbolic representations as well should be treated as the social units which are subject to this constitutional dynamics. This broad concept of the social unit has an interesting consequence for the design of multi-agent systems. The agents of the multi-agent system are modelled to represent social units and not only social actors in the traditional sense. Thus, the multi-agent system allows to represent the social agency of social and symbolic structures.

Modelling and Exploring hybrid constellations

8.1

Multi-agent systems consist of individuals — the agents — with capabilities of their own and with a certain degree of autonomy in initiating and conducting activities. Through social processes of interaction and exchange these individual agents shall work together to perform common tasks or to get support from other agents for their individual tasks. Multi-agent systems, thus, are social systems by definition. A large part of the communication and interaction protocols and of the coordination mechanisms in multi-agent systems are adopted from human societies, since human societies provide a rich reservoir of solutions for the problem of coordination and cooperation between autonomous entities. For this reason it is promising to combine the science of human societies — sociology — with the research on multi-agent systems. However, to use social mechanisms of human societies as models for multi-agent coordination has implications that go far beyond constructing more powerful or more sophisticated multi-agent systems. It implies as a next step the construction of common social systems of human and artificial agents, that is, hybrid settings with mechanisms of social interaction and exchange which are commonly used by humans and by artificial agents.

8.2

The INKA project (Meister et al., in this special section) explores such a hybrid constellation. The project's domain of application is negotiations on shift exchanges in a hospital. Negotiating and trading single working shifts is a common practice in hospitals. It is used by employees to realise individual leisure-time interests that are not considered in the official rosters. Exchanging shifts is an informal process based on experiences the employees make with one another. Over time these experiences become patterned expectations. One comes to know the other for instance as a family-oriented person to whom it is difficult to trade day shifts or as an agreement-oriented character whose willingness to compromise is high. The INKA researchers conceptualise these patterned expectations as practical roles. The different expectation patterns built on the basis of empirical observation are conceptualised as social types. In a first step the INKA project has developed a multi-agent system where the different agents act in accordance to the different practical roles of these social types. Provided with certain leisure-time interests which are in conflict with the official roster, these agents are capable of trading shifts on the basis of taking into account their own social type as well as the social type of their negotiation partners. In a second step this multi-agent system was introduced into the real-world setting of human shift negotiations by developing a common interface that allows shift negotiations between humans, between humans and agents, and between agents. The performance of the different constellations was tested in a so called interactivity experiment. The experiment was designed as role-play scenario within a laboratory setting. Employees from a hospital who volunteered as probands where instructed to act in the role of the different social types and to negotiate shifts by means of the computer interface not knowing whether their respective negotiation partners were humans or agents. The results for human-human and human-agent negotiations stemming from this experiment are supplemented by simulation runs of agent-agent negotiations.

8.3

It is not realistic to assume that such a negotiation system will be used in hospitals in the foreseeable future — if ever. So what is the use of developing and exploring it? The answer to this question goes back to the answer to the question of why multi-agent systems are useful at all. The pioneers of distributed computing have argued that distributed problem-solving is preferable to centralised computing in situations where problems occur locally, where the relevant information has to be gathered from different spaces, and where problem-solving competencies are distributed over several units. The rationale of distributed problem-solving is that local problems and local information should be dealt with by local problem-solvers and that the overall solution should result from the interaction of the local problem-solvers.

8.4

With respect to their formal structure, organisations basically are centralised problem-solvers. However, much empirical research has shown that the solutions of many of the day-to-day problems within organisations follow the pattern of distributed problem-solving. The problems occur locally and are solved by local knowledge and practices leading to informal structures which supplement an organisations formal rules and procedures. Negotiating working shifts is such a mechanism of distributed problem-solving that deals with local problems (leisure-time interests) by employing local knowledge (expectations about possible exchange partners) and that is employed to overcome deficiencies of the official roster as the centralised problem-solving mechanism of the formal organization. The aim of the INKA project is to explore hybrid settings in which human actors and artificial agents work together as distributed problem-solvers within a common frame of reference of social interaction. The aim is not to develop a marketable product. Thus the domain of application is well chosen. Beyond describing the agent architecture and the common frame of reference of social interaction the paper deals with the methodological question of how to evaluate the performance of hybrid settings in comparison with pure human-human or pure agent-agent settings.

Acknowledgements

Today the members of the Socionics program can look back on six years of successfully treading together the thorny path of interdisciplinary cooperation. During these years many ideas and concepts were exchanged and elaborated, many different models were constructed, prototypes were developed and tested, conferences and workshops held, networks and friendships established, and hundreds of articles published and several books[4] edited. Last but not least a collection of five selected research papers has been prepared for publication in this special section of the *JASSS*. On behalf of all contributing colleagues we want to thank the anonymous referees for their comments and suggestions which helped to improve the original draft versions of these papers. Our special thanks go to Nigel Gilbert who invited us to publish these reports in the *JASSS* and who encouraged and supported us in the laborious process of preparing this special section. We also want to seize the opportunity to thank the group of advisors from the DFG who evaluated our bi-annual applications and research reports and who, together with Helga Hoppe and Gerit Sonntag from the administrative staff, accompanied us as critics and solicitors on all issues of the Socionics program.

Notes

[1] Hewitt and Gasser, in collaboration with sociologists like Gerson and Star, were among the first computer scientists to initiate a dialogue with sociologists and to discuss the question of how DAI research could be build on sociological foundations (Gasser et al. 1989, Gasser 1991, Hewitt 1991). For details cf.Strübing 1998.

[2] The initiators' group consisted of two computer scientists, Wilfried Brauer and Hans Jürgen Müller, and two sociologists, Thomas Malsch (Coordinator) and Werner Rammert.

[3] Cf. the research initiative "Complex Systems" within the IST-Workprogram 2003-2004 and the EU funded "Complex Systems Network of Excellence".

[4] Malsch 1998; Kron 2002; Rammert & Schulz-Schaeffer 2002; v. Lüde, Moldt & Valk 2003; Florian & Hillebrandt 2004; Lindemann, Moldt & Paolucci et al. 2004; Fischer, Florian & Malsch 2005; Schmitt, Florian & Hillebrandt 2006.

References

AHRWEILER, P and Gilbert, N (Eds.) (1998) *Computer Simulations in Science and Technology Studies*. Berlin: Springer.

ARTIKIS, A and Pitt, J (2001) "A Formal Model of Open Agent Societies". In: Müller, J (Ed.) *Proceedings of the Fifth International Conference on Autonomous Agents (AA)*, New York: ACM Order Dep.. pp. 192-193.

BONABEAU, E, Dorigo, M and Theraulaz, G (1999) *Swarm Intelligence: From Natural to Artificial*

Systems. New York: Oxford University Press.

BOND, A H and Gasser, L (Eds.) (1988) *Readings in Distributed Artificial Intelligence*. San Mateo, Ca.: Morgan Kaufmann Publishers.

BRAFMAN, R I and Tennenholtz, M (1996) On Partially Controlled Multi-agent Systems. *Journal of Artificial Intelligence Research*, 4. pp. 477-507.

CHIA, M H, Neiman, D E and Lesser, V R (1998) Poaching and distraction in asynchronous agent activities. *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98)*, Paris. pp. 88-95.

CONTE, R, Falcone, R and Sartor, G (Eds.) (1999) Special Issue: Agents and Norms. *Artificial Intelligence and Law 7 (1).*

DAVIDSSON, P (2000) "Emergent Societies of Information Agents". In Klusch, M and Kerschberg, L (Eds.) *Cooperative Information Agents IV: the future of information agents in cyberspace*, Berlin: Springer. pp. 143-153.

DAVIDSSON, P (2001) "Categories of Artificial Societies". In Ominici, A, Petta, P and Tolksdorf, R (Eds) *Engineering Societies in the Agent World II: Second International Workshop (ESAW)*, Berlin: Springer. pp. 1-9.

DAVIS, R and Smith, R G (1983) *Negotiation as a Metaphor for Distributed Problem Solving. Artificial Intelligence*, 20/1983. pp. 63-109.

DELLAROCAS, C and Klein, M. (1999) Civil Agent Societies: Tools for Inventing Open Agent-Mediated Electronic Marketplaces. *Proceedings Workshop in Agent-Mediated Electronic Commerce (co-located with IJCAI'99)*. Stockholm: ACM Press. pp. 24-39.

DELLAROCAS, C and Klein, M. (2000) An Experimental Evaluation of Domain-Independent Fault Handling Services in Open Multi Agent Systems. *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS)*. Los Alamitos: IEEE Computer Soc.. pp. 95-102.

DORAN, J and Gilbert, N (1994) "Simulating Societies: An Introduction". In Gilbert N and Doran J (Eds.) *Simulating Societies. The Computer Simulation of Social Phenomena*, London: UCL-Press. pp. 1-18.

EDMONDS, B and Dautenhahn, K (Eds.) (2001) Special Section: Starting from Society: the Application of Social Analogies to Computational Systems. *Journal of Artificial Societies and Social Simulation* 4(1) http://jasss.soc.surrey.ac.uk/4/1/0.html.

FISCHER, K, Florian, M and Malsch, T (Eds.) (2005) *Socionics: Its Contributions to the Scalability of Complex Social Systems*. Berlin: Springer.

FLORIAN, M and Hillebrandt, F (2004) (Eds.) *Adaption und Lernen von und in Organisationen. Beiträge aus der Sozionik*. Wiesbaden: VS Verlag.

GASSER, L (1991): Social Conceptions of Knowledge and Action: DAI Foundations and Open Systems Semantics. *Artificial Intelligence*, 47/1991. pp. 107-138.

GASSER, L, Rouquette, N F, Hill, R H and Lieb, J (1989) "Representing and Using Organizational Knowledge in Distributed AI Systems". In Gasser L and Huhns N M (Eds.) *Distributed Artificial Intelligence, Volume II*, London: Pittman, San Mateo, Ca.: Morgan Kaufmann. pp. 55-87.

GUSTAVSSON, R and Fredriksson, M (2001) "Coordination and Control in Computational Ecosystems: A Vision of the Future". In Omicini, A, Klusch, M, Zambonelli, F and Tolksdorf, R (Eds.) *Coordination of Internet Agents: Models, Technologies, and Applications*, Berlin: Springer. pp. 443-469.

HALES, D (2001) *Tag Based Co-operation in Artificial Societies*. unpublished PhD-Thesis, Department of Computer Science, University of Essex http://cfpm.org/~david/thesis.

HEWITT, C E (1977) Viewing Control Structures as Patterns of Passing Messages. *Artificial Intelligence*, 8/1977. pp. 323-364.

HEWITT, C E (1991) Open Information Systems Semantics for Distributed Artificial Intelligence. *Artificial Intelligence*, 47/1991. pp. 79-106.

HUHNS, M N and Singh, M P (1998) *Readings in Agents.* San Francisco: Morgan Kaufmann.

JENNINGS, N R (1993) Commitments and Conventions. The Foundation of Coordination in Multi-agent Systems. *Knowledge Engineering Review*, 8(3). pp. 223-250.

KEPHARD, J and Chess, D M (2003) The Vision of Autonomic Computing. *IEEE Computer*, 36(1). pp. 41-50.

KLUSCH, M (2001) Information Agent Technology for the Internet: A survey. *Journal on Data & Knowledge Engineering*, 36 (3). pp. 337-372.

KRON, T (Ed.) (2002) *Luhmann modelliert: Sozionische Ansätze zur Simulation von Kommunikationssystemen*, Opladen: Leske & Budrich.

LINDEMANN, G, Moldt, D, Paolucci, M and Yu, B (Eds.) (2004) *International Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA'02). Post Proceedings*. Berlin: Springer.

LUCK, M, McBurney, P and Preist, C (2003) "Agent Technology: Enabling Next Generation Computing (A Roadmap for Agent Based Computing)", Online-Publikation AgentLink http://www.ecs.soton.ac.uk/~mml/papers/al2roadmap.pdf.

MALSCH, T (2001): Naming the Unnamable: Socionics or the Sociological Turn of-to Distributed Artificial Intelligence. *Autonomous Agents and Multi-agent Systems*, 4(3). pp. 155-186.

MALSCH, T (Ed.) (1998) *Sozionik: soziologische Ansichten über künstliche Sozialität*. Berlin: Edituin Sigma.

MALSCH, T, Florian, M, Jonas, M, Schulz-Schaeffer, I (1996) Sozionik: Expeditionen ins Grenzgebiert zwischen Soziologie und Künstlicher Intelligenz. *Künstliche Intelligenz*, 10(2). pp. 6-12.

MALSCH, T (1997) Die Provokation der "Artificial Societies". Warum die Soziologie sich mit den Sozialmetaphern der Verteilten Künstlichen Intelligenz beschäftigen sollte. *Zeitschrift für Soziologie,* 26(1). pp. 3-22.

MAMEI, M and Zambonelli, F (2004) "Self-Organization in Multi Agent systems: A Middleware Approach". In Di Marzo Serugendo, G, Karageorgos, A, Rana, O F and Zambonelli F (Eds.) *Engineering Self-Organising Systems, Nature-Inspired Approaches to Software Engineering*. Lecture Notes in Computer Science 2977, New York: Springer. pp. 233-248.

MÜLLER, H J (1993) "Introduction". In Müller H J (Ed.) (1993) *Verteilte Künstliche Intelligenz. Methoden und Anwendungen*, Mannheim: BI-Wissenschaftsverlag. pp. 9-21.

MÜLLER, H J, Malsch, T and Schulz-Schaeffer, I (1998) Socionics: Introduction and Potential. *Journal of Artificial Societies and Social Simulation* 1(3) http://jasss.soc.surrey.ac.uk/1/3/5.html.

OMICINI, A, Tolksdorf, R and Zambonelli, F (Eds.) (2000) *Engineering Societies in the Agent World*. Berlin: Springer.

OSSOWSKI, S (1999): *Coordination in Artificial Agent Societies: Social Structures and its Implications for Autonomous Problem-solving Agents*. Berlin: Springer.

PANZARASA, P and Jennings, N R (2001) The Organisation of Sociality: A Manifesto for a New Science of Multi-agent Systems. *Proceedings of the 10th European Workshop on Multi-agent systems (MAAMAW)* http://www.ecs.soton.ac.uk/~nrj/download-files/maamaw01.pdf.

PAPADIMITRIOU, C H (2001): Algorithms, Games, and the Internet. *Proceedings 33rd Ann. ACM Symposium Theory of Computing*. New York: ACM Press. pp. 749-753.

PRIETULA, M, Carley, K and Gasser, L (eds.) (1998) *Simulating Organizations: Computational Models of Institutions and Groups*. Cambridge, MA: The MIT Press.

RAMMERT, W and Schulz-Schaeffer, I (2002) *Können Maschinen handeln? Soziologische Beiträge zum Verhältnis von Mensch und Technik*. Frankfurt/Main: Campus.

SCHMITT, M, Florian, M and Hillebrandt, F (Eds.) (2006) *Reflexive soziale Mechanismen. Von soziologischen Erklärungen zu sozionischen Modellen*. Wiesbaden: VS Verlag.

STRÜBING, J (1998) Bridging the Gap: On the Collaboration between Symbolic Interactionism and Distributed Artificial Intelligence in the Field of Multi-Agent Systems Research. *Symbolic Interaction*, 21. pp. 441-464.

VERHAGEN, H J E (2000) *Norm Autonomous Agents*. unpublished PhD-Thesis. Stockholm University/Royal Institute of Technology.

VON LÜDE, R, Moldt, D, and Valk, R (2003) *Sozionik - Modellierung soziologischer Theorie*. Münster: LIT.

WAHLSTER, W and Weyrich, C (Eds.) (2002) Forschen für die Internet-Gesellschaft: Trends, Technologie, Anwendungen. Ergebnisse einer gemeinsamen Initiative des Bundesverbands der Deutschen Industrie und der Fraunhofer Gesellschaft http://w4.siemens.de/ct/internet-trends.

WELLMAN, M P and Wurman, P R (1998) Market-Aware Agents for a Multi-Agent World. *Robotics and Autonomous Systems*, 24. pp. 115-125.

ZAMBONELLI, F and Parunak, V (2003) Towards a Paradigm Change in Computer Science and Software Engineering: A Synthesis. *The Knowledge Engineering Review*, 18(4). pp. 329-342.

Return to Contents of this issue