LEARNING FROM AND ABOUT OTHER AGENTS IN TERMS OF SOCIAL METAPHORS

Michael Schillo, Petra Funk
DFKI GmbH, Im Stadtwald, Geb. 36
D-66123 Saarbrücken, Germany
e-mail: {schillo, funk}@dfki.de
Tel.: +49 681 302 2464
Fax: +49 681 302 2235

Abstract

We present work that has been conducted in a sociological and psychological context. Our aim is to establish a mechanism that enables agents to cope with environments that contain both selfish and co-operative entities, where the mixture and the behaviour of these entities is previously unknown to all agents. We achieve this by enabling agents to evaluate *trust* in others, based upon the observations they gather. Using trust, they are able to request observations from others and make use of this possibly manipulated data, therefore enlarging their data base on behaviour of others. Our approach results in significantly faster and better behaviour adaptation. We demonstrate the improvement in performance of agents using trust compared to the performance of others that just use their own observations.

1. Motivation

Multi-agent systems typically are of a complex and dynamic nature. Therefore, for a designer of such systems it is inherently impossible to pre-specify all behaviours and activities of a single agent or of a collection of concurrently active agents. Learning and adaptation offer elegant concepts and methods to solve these problems. Until recently, most multi-agent systems implicitly incorporated benevolence and co-operation as key features. With open systems and virtual co-habited worlds, such as electronic market places or other internet forums, multi-agent systems evolve from pure collections of co-operating units to artificial societies enriched with social structures and attitudes. Socially intelligent agents, and their ability to cope with egoistic or lying actors are new challenges arising within this development (Dautenhahn et al. 1997). Agents may find themselves confronted with deception and fraud; being able to reason about other's intentions before committing to actions, becomes essential to survival in an artificial social setting.

As a solution to this problem, we present a conceptualisation of trust, which is established by an observation and communication process. Agents start out with no knowledge about other members in the society and then modify their model on trustworthiness of others according to observation and testimony from others. During several interaction activities, the model about other members of the society is refined and used to judge their reliability to commitments to announced activities. Thus, agents are able to approximate other agents intentions, when offered co-operation.

2. Social metaphors – Socionics

In this section we will explain why and to what extend we use social metaphors for our approach, and how they relate to learning and adaptation processes.

One attempt to introduce terminology from social sciences is the work of Bazzan et al. (Bazzan et al. 1997). The authors tried to create a society of more successful agents by providing the concepts of altruism, egoism and using altruists with moral sentiments to support themselves. Another important influence to this work is the newly created discipline of *socionics*, a recent effort by the German sociologist Thomas Malsch. In this effort, researchers intend to merge research from the two disciplines *distributed artificial intelligence* and *sociology* in order to create new insights to a number of problems that are common to both disciplines, e.g. the problems of distributedness, heterogeneity, emergence and problem solving (Malsch et al. 1996)[1].

[1] *The scientific process is reflected by a funding program of the Deutsche Forschungsgemeinschaft (the German equivalent of the NSF) that supports research on and modelling of artificial societies.*

One of the first assumptions we rejected, was the inherent benevolence assumption in multi-agent systems, a step already taken by Bazzan et al. In contrast to their work, we did not choose to replace this assumption by a set of fixed behaviours, but by fuzzy social attitudes. In our work, these attitudes can range from *egoistic* to *altruistic*, from *honest* to *dishonest*. As described before, our aim was to model artificial societies where the agents did *not* know about the attitudes of their opponents. So they must learn by observation in order to create reliable models.

Also, we wanted enable agents to adapt very quickly to a new environment. Thus it is essential for them to learn from other agents that have encountered a similar situation before. However, intentions and motives can differ from agent to agent and an agent cannot hope that every other agent wants to co-operate. It may be the case that other agents will use their influence (*power*), when asked by others for their observations, to try to manipulate them by communicating incorrect data (*lying* and *deception*).

Having outlined the problem with these terms, the solution is provided by social sciences in a rather straightforward manner: an agent has to learn which witnesses it can *trust* in order to provide it with meaningful data and to evaluate how far it can trust these witnesses. Note that this is not a problem about how accurate the data of the witnesses is, but of how willing the witness is to communicate it, be it accurate or not. Thus, in this case we are talking about the trustworthiness of the witness and not about the *reliability* of the data of the witness. Dealing with such brittle testimonies is a hard learning task that will end (if successful) in a behaviour adaptation, so that the agent learns by observation and communication which other agents are co-operative partners and will therefore know which agents to seek and which to avoid in interactions.

The social science solution alone will not solve the problem in terms of algorithms and computer programs. Therefore we developed a computational model of trust and a new and superior way to calculate it than it has been done before ((Marsh 1994), (Beth et al. 1994)). We achieve this with the aid of a graph-like data structure that stores the agent's observations and the testimonies that the witnesses communicated. Computing the model according to this data is then a process based on probabilistic assumptions which can be found in more detail in (Schillo und Funk 1999). Due to the complexity of the object of sociological research we had to use definitions of terms that may not always be the broadly accepted definitions in the field of sociology. However, using this terminology enables us to bridge the gap between sociology and DAI and helps us to describe our system in a more elegant way.

3. Using the Trust Model

The evolution of multi-agent systems from socially neutral to socially rich artificial societies offers new application domains, which are hosted by open systems such as the internet. Electronic commerce is a broad application domain comprising severe traps to blindly trusting agents. The commonly suggested solutions of Trusted Third Parties as global authorisers requires mutual agreement by all designers of services and may therefore be hard to achieve. Other exogenous control methods, as have been described in (Armstrong and Durfee 1998), the authors there propose library mechanism selling information to one another. Our approach helps each single agent to establish a model of trustworthiness of other agents (or net services). With only few iterations, as will be shown in the results section, agents learn whom to trust and whom to exclude from future interactions.

To verify our claims and the underlying implementation, we developed a variant of the well-know prisoner's dilemma ((Luce and Raiffa 1957), (Axelrod 1984)), which we call *disclosed prisoner's dilemma with partner selection*. We have introduced this variation of the game elsewhere (Schillo and Funk 1998); here we outline a slightly modified version of it. The partner selection offers agents the possibility to use their model about the trustworthiness of other agents, which has been gained through observation and testimonies. Before starting the actual game, agents pay a stake, in order to participate in the next round. Then, in a contract net-like style, an agent announces its intention to play and also, how it intends to play. The other agents then can decide to make a bid, where they themselves announce their willingness to play and also, how they intend to react to the announced move. Each bidder will base its announcement on the information gathered about the current manager of the negotiation. If a bidder has not sufficient information about this agent, it may ask some other agents about their observations on the reliability of the agent in question. It will then have to evaluate the information gathered by the existing trust model and use the information for refinement of the model. After receiving all bids, the manager evaluates these in the same manner: It checks the trustworthiness of the bidders according to its own trust net. Like in the bidders case, the manger may not have enough information to evaluate the bid. It may then engage in communication with other agents in order to gather information from them about the bidder. When the manager has chosen its partner for this round, another agent becomes manager and announces its intentions to play. It is important to note here that during this partner selection phase, agents are free lie about their observations.

It may e.g. happen, that the manager asks an agent, who would like to play with this manager, for information about another bidder. Since the agent would rather play itself, it may lie about the trustworthiness of the agent in question in order to get the award.

4. Behaviour and Adaptation

Our agents behave according to a probabilistic model. Depending two parameters they are more or less likely to behave altruistic/egoistic or honest/dishonest. The modelling of the behaviour consists in the task to approximate this parameter for each agent. This approximation is achieved by gathering data on past behaviour and then evaluate averages. The crucial point here is that our agent's do not exclusively rely on their own observations, but they also take into account the testimonies of witness agents. As these may be lying, every agent that receives data has to evaluate the witnesses behaviour. If there is not enough data on these witnesses, the agent will recursively apply the mechanism described above to gather data on them, again using both, observations and testimonies.

The behaviour adaption of the agents consists in the choice of interaction partners. The proposed agents will not change their intended behaviour as specified by their honesty and altruism parameter. Instead they will adapt their behaviour in the sense that they decide not to interact with agents that behave in an undesired manner and rather try to interact with agents that play e.g. co-operation more often. This is a different approach to the one taken by Sen et al. ((Sen 1996), (Biswas et al. 1999)). Here, the authors enable agents to evaluate a best response to agents that they need to interact with, even if the strategies of these agents are unknown. In the application domains targeted by our work (see next section), agents are able to reject an interaction. Therefore we must enable agents to make a qualified decision with whom to play.

5. Application Domains

In this section we show the connection between the described setting and real world applications, where trust and finding out whom to trust is an essential goal. We choose the electronic commerce application to demonstrate practical relevance.

In both settings, in electronic commerce as well as in our scenario, agents encounter other agents autonomously interacting in order maximise their performance. They may achieve this by co-operating with their contract partners in the long run, or try to make "fast cash" and exploit others. Additionally, agents can offer contracts and certain behaviour that they do not intend to exhibit at commitment time. Electronic commerce is unbound by national frontiers and therefore free from national authorities. In traditional trading, however, these authorities guaranteed the fulfilment of agreements or punishment otherwise by their power. This means of stability is not available in electronic commerce (and for many business partners it will not be for some time). In our setting, the agents will not be punished if they deceive their game partners, but we enable agents to track very fast which agents are behaving in a deceitful way so that they can avoid them. To model the conflict between behaving co-operative or deceitful and the respective outcomes, we use the prisoner's dilemma with the addition that agents need to pay before they are allowed to join a round of the game. This results in a loss of score if they do not find someone that will join them in a game. Thus, there is an indirect punishment to malicious behaviour, which can be described as *virtual peer pressure*.

The idea of the peer pressure is based on the natural assumption that there is some communication in the world about the players in the market, e.g. press, personal communication, etc. Again this communication is in both settings not completely objective (to avoid the somewhat ambiguous term *trustworthy*). In our experiments, agents will be lying to some extend. As they do not want to be found making up data, they will try to bias the information they communicate by leaving out the data they have observed that does not match to their intentions. We assume that they are motivated to make their competitor look *not* trustworthy and *not* suited for games with high pay-off in order to discourage players to choose one of them.

Electronic commerce and our setting have many features in common and both are dangerous places for agents relying on benevolent contractors. Agents which succeed in our setting will be interesting to investigate, when looking at the real world application. Additionally, we currently envision two other applications of the proposed mechanism. We believe that in the future mechanism like „Cookies" will be of more importance in the Internet world. The demonstrated solution will be an easy to implement way to find out say, which host you can trust to not abuse the information of the Cookie set on your machine (e. g. for using personal data). More serious is a similar application where our approach is just as applicable: the problem of differentiating between friendly and malicious hosts or migrating programs (see e.g. (Sander and Tschudin 1998) or (Vigna 1998)).

6. Results

Using social metaphors to describe learning in multi-agent systems does not only provide elegant models, it also delivers significant results in terms of performance increase. In this section we present a number of results that have been produced by simulation experiments using the *disclosed prisoner's dilemma with partner selection* (Schillo and Funk 1998) and a society of 40 agents in 20 different configurations that compete for the best pay-off. As mentioned before, our agents are characterised by two parameters: the probability for honesty/dishonesty and for altruism/egoism. We chose four different values for honesty (0.01, 0.33, 0.66 and 0.99) and five values for altruism (0.01, 0.25, 0.50, 0.75 and 0.99). We included every possible combination of these values, which results in the 20 different agent strategies (i.e. configurations). So this society is highly heterogeneous.

As platform for our simulations we chose the *Social Interaction Framework* (*SIF*) that has been developed at the DFKI[2]. The code for the system and the implementation of the trust evaluation component (TrustNet) is written in Java™ and is available via ftp. For analysis we used the data of seven simulations during 300 hours of computation time, distributed over six Linux PCs and a SPARC Ultra 2. From the convergence of the data, we can conclude that this set of data is representative for the setting and its configuration.

In the simulations, we had two groups of agents, an experimental group that used the model of trust to evaluate possible game partners and a control group that only used its own observations for this evaluation. The produced data shows that using the TrustNet is of benefit to the agents.

First, we take a look at some example data to get a flavour of the improvement of the performance of these agents. For this first analysis we chose agents that had a high degree of egoism and deception (Figure 1) and agents with high altruism and high honesty (Figure 2). These two configurations can be looked at as the two extremes in an application: a malicious host or contract partner (with altruism 0.25 and honesty 0.33: configuration 0.25 0.33) and a friendly mobile agent or a electronic commerce customer (with altruism 0.99 and honesty 0.99: configuration 0.99 0.99). As performance we measure the accumulated pay-off in an iterated game of *disclosed prisoner's dilemma with partner selection* (a slighlty modified *prisoner's dilemma*). Figure 1 shows that in the short term the performance of the

[2] Further information about SIF can be found at http://www.dfki.de/~sif/

agents using the TrustNet (the experimental group E 0.25 0.33) is superior to the corresponding control group (C 0.25 0.33).

The performance of the experimental group is flattening with the growing quality of the models of all agents and their growing intention not to play with agents from this group. The other figure shows analogous results for a group of altruists. The experimental group (E 0.99 0.99) is by far more successful than the control group (C 0.99 0.99) and after 25 games the surplus in performance is still growing. This surplus is due to the fact that the altruists from the experimental group do not let themselves be deceived as many times as the control group agents, which is shown by the graphs. The experimental group of altruists does not have any setback in score at all, whereas the control groups suffer from at least two such setbacks (note that these graphs are averages on the conducted simulations).

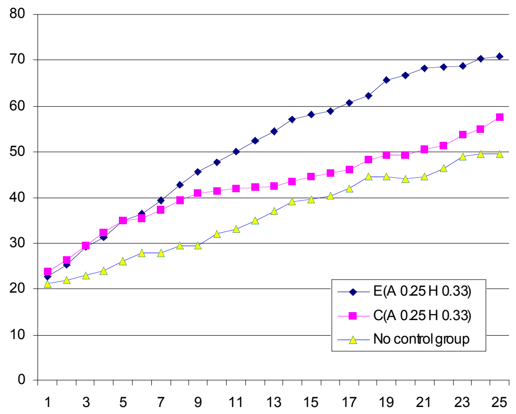

Figure Egoist performance

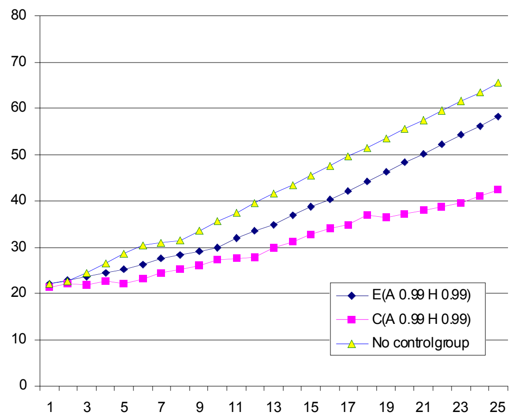

Figure 2 Altruist performance

To some extend, our groups seem to behave sub-optimal: the altruists in Figure 2 do not have a better performance than the egoists in Figure 1, which is not desired. Further simulations with societies of agents which all used the TrustNet (no control group) have shown that the reason to this is the social intelligence of the society as a whole. If we have only socially intelligent agents, the performance of the egoists decreases (as they find less agents to deceive) and the performance of the altruists increases (as they find more agents which trust them to be altruists). This improved setting and the corresponding performances are shown by the graphs annotated with ‘no control group’ in the two figures. These graphs show that the altruists can perform better than the egoists, if they use the TrustNet and that the experimental group is better than the control group. However, agents can rely on the social intelligence of the other members of the society, therefore we continue simulating mixed societies.

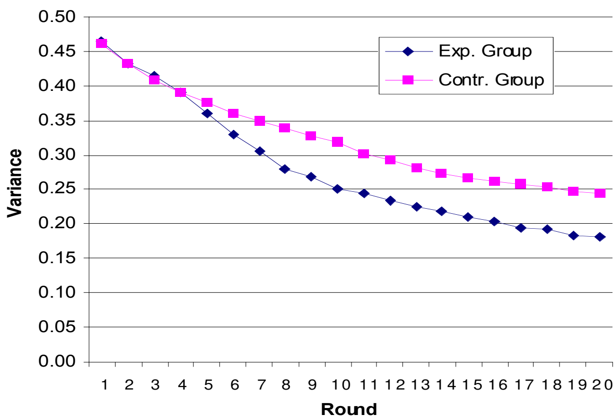

Figure 3 The variance of the error in the agent's models.

This gain in performance is due to the improved model of others, as the agents have learned more and faster about the behaviour of others. In order to measure this we calculated for each group of agents the error in the learned approximation to the actual behaviour of any other member of the society. The average variance of this error is shown in Figure 3. This shows that agents using the TrustNet have already after round ten a model with 20% less variance. This results in the need for the control group to learn through eight more rounds to reach a similar quality of their model.

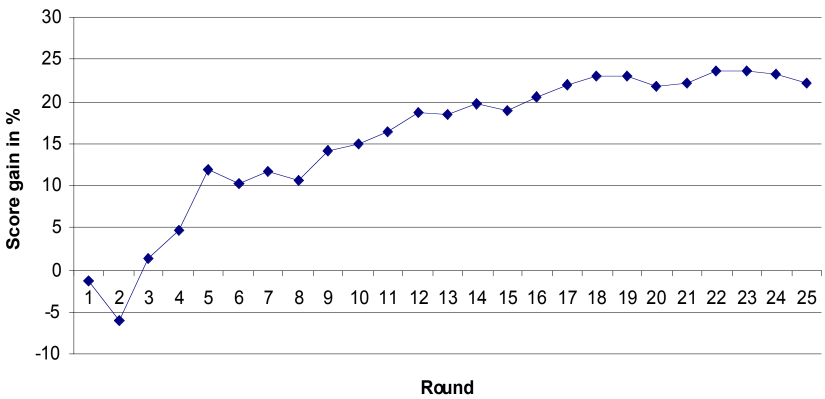

Figure 4 The gain of performance when using trust.

This advantage is of major importance to the overall performance of the agents in the society. During the first 25 rounds in a society of unknown configuration, the performance of the agents develops a significant improvement compared to the agents that just use their observations. The total gain of score of the agents in the experimental group is more than 20% higher and reaches a level of plus 10% after five rounds.

7. Conclusion

7.1. Summary

With our concept of trust and its formalisation we have provided a mechanism that allows for better adaptation to the behaviour of previously not assessed entities. This is managed by enabling agents to use data communicated by others, even if these witnesses can be assumed to be deceptive about their knowledge and are manipulating this data. With this feature, the amount of observational data is enlarged, which allows for a quicker and better approximation of features to an extend that justifies its application in settings like electronic commerce and the mobile agents problem.

We analysed material that shows how much the performance of agents can improve when our approach is used. The effect is that agents can learn models from scratch almost twice as fast as other agents that only use their own observations, while still reaching the same or better accuracy. They achieve this even when confronted with witnesses that communicate manipulated data. Furthermore, we enable co-operative agents to form groups and play among themselves and thus, when having worked out who belongs to this group, forming more and more stable groups of agents that profit from mutual support.

During the whole process of this work, we have used social metaphors as a language for describing interaction in multi-agent systems. We consider this a natural effect, as we assume a close analogy between multi-agent systems (i.e. artificial societies) and human societies.

7.2. Future work

Trust is a mechanism for creating trustworthy channels for communication and can by nature be used for communicating a range of features important to multi-agent systems. We intend to explore further applications, the use of other features than altruism and evaluate the changes in performance.

While this paper is being compiled, we are conducting experiments in order to pin down under which circumstances using the TrustNet will increase performance and to what extent. We intend to determine how quickly agents can adapt by using the TrustNet if they enter a society of agents, who already have learned about each other's attitudes. We are also interested in researching the effects of changes of the behaviour over time and especially in more complex forms of lying that are guided by strategies and how agents can adapt to this.

Varying the ratio of egoists and altruists leads to different distribution curves on the agent society structure. Researching the effects of the TrustNet in such different society structures is another interesting challenge.

Acknowledgements

We would like to thank the workshop reviewers for their detailed and helpful questions and suggestions on earlier versions of this paper.

Bibliography

(Armstrong and Durfee 1998) Armstrong, A. A. and Durfee, E. *Mixing and Memory: Emergent Cooperation in an Information Marketplace*. In: Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS 98), 1998.

(Axelrod 1984) Axelrod, R. *The Evolution of Cooperation*. New York: Basic Books, 1984.

(Bazzan et al. 1997) Bazzan, A. L. C., Bordini, R. H. und Campbell, J. A. Agents with Moral Sentiments in an Iterated Prisoner's Dilemma Exercise. In:(Dautenhahn et al. 1997).

(Beth et al. 1994) Beth, T., Borcherding, M. und Klein, B. Valuation of Trust in Open Networks. In: *Computer Security-ESORICS 94 – Third European Symposium on research in Computer Security*. Brighton, UK, November, 1994.

(Castelfranchi et al. 1999) Cristiano Castelfranchi, Yao-Hua Tan, Rino Falcone, Babak Sadighi Firozabadi. *Deception, Fraud and Trust in Agent Societies*. Proceedings of the workshop at the Autonomous Agents Conference. National Research Council, Institute of Psychology, Rome, Italy, 1999.

(Dautenhahn et al. 1997) Dautenhahn, K., Masthoff, J. und Numaoka, C. *Socially Intelligent Agents*. Papers from the 1997 AAAI Fall Symposium, November 8-10, Cambridge, Massachusetts, Technical Report FS-97-02, 1997.

(Luce and Raiffa 1957) Luce, R. and Raiffa, L. *Games and Decisions*. Wiley, New York, 1957.

(Malsch et al. 1996) Malsch, T. et al. *Expeditionen ins Grenzgebiet zwischen Soziologie und Künstlicher Intelligenz.* In: Künstliche Intelligenz, 2:96, S. 6-12, in german, 1996.

(Marsh 1994) Marsh, S. P. *Formalising Trust as a Computational Concept.* Phd Thesis, Department of Computing Science and Mathematics, University of Stirling, 1994.

(Sander and Tschudin 1998) Sander, T. and Tschudin, C. F. Fritz. *Protecting Mobile Agents against Malicious Hosts*. In: (Vigna 1998), S. 44 ff., 1998.

(Sen 1996) Sandip Sen. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In *Proceedings of the Second International Conference on Multiagent Systems*, AAAI Press, Menlo Park, CA, 1996.

(Biswas et al. 1999) Anish Biswas, Sandip Sen, Sandip Debnath. Limiting Deception in Social Agent-Group. In (Castelfranchi et al. 1999).

(Schillo and Funk 1998) Schillo, M. and Funk, P. Spontane Gruppenbildung in künstlichen Gesellschaften. In: *Proceedings des Workshops Sozionik der 22. Jahrestagung für Künstliche Intelligenz.* In German, 1998.

(Schillo und Funk 1999) Michael Schillo und Petra Funk. Who can You Trust: Dealing with Deception. In Cristiano Castelfranchi, Yao-Hua Tan, Rino Falcone, Babak Sadighi Firozabadi. *Deception, Fraud and Trust in Agent Societies*. Proceedings of the workshop at the Autonomous Agents Conference. National Research Council, Institute of Psychology, Rome, Italy, 1999.

(Vigna 1998) Vigna, G. *Mobile Agents and Security*. Springer-Verlag, Berlin, 1998