Argumentation and multi-agent decision making

Simon Parsons and Nick R. Jennings

Department of Electronic Engineering,
Queen Mary and Westfield College,
University of London,
London E1 4NS,
United Kingdom.

{S.D.Parsons,N.R.Jennings}@qmw.ac.uk

Abstract

This paper summarises our on-going work on mixed-initiative decision making which extends both classical decision theory and a symbolic theory of decision making based on argumentation to a multi-agent domain.

Introduction

One focus of our work at Queen Mary and Westfield College is the development of multi-agent systems which deal with real world problems, an example being the diagnosis of faults in electricity distribution networks (Jennings *et al.* 1996). These systems are mixed-initiative in the sense that they depend upon interactions between agents—no single agent has sufficient skills or resources to carry out the tasks which the multi-agent system as a whole is faced with. Because the systems are built to operate in the real world, the agents are forced to deal with the usual problems of incomplete and uncertain information, and increasingly we are turning to the use of techniques from decision theory, both classical and non-standard, in order to ensure that our agents make sound decisions.

This position paper briefly describes three lines of investigation within our programme which can be thought of as research into mixed-initiative, multi-agent decision making in the sense described above. These are:

- Agents as lone decision makers. Any autonomous agent in an uncertain world is a resource bounded entity which must make decisions under uncertainty. Belief/Desire/Intention (BDI) models have been proposed as a mechanism for building such agents but they lack a fine-grained representation of uncertainty. We are seeking to extend BDI models with more sophisticated models of uncertainty.

- Socially aware agents. Building agents which make use of classical decision theory leads to groups of selfish utility maximisers. Since it appears that, in certain circumstances, societies of agents may benefit from more altruistic agents, we are investigating ways of making agents more socially aware within the context of classical decision theory.

- Agents as information sharers. In order for groups of agents which do not have perfect information to reach joint decisions, they need to share information. We are investigating how this may be done through dialogue between agents, and how dialogues may be guided by the use of argumentation, a mechanism which also provides a symbolic model for agent decision making.

In the remainder of the paper we consider each of these three areas in (slightly) more detail along with their relation to mixed-initiative decision making and work on agents in general, and also mention some related work in which we are involved.

Lone decision makers

Autonomous agents need a mechanism of making decisions about what to do. A popular way of achieving this is to ascribe agents mental attitudes, and to talk about their beliefs (what they believe), their desires (what they want) and their intentions (what they try to bring about). The resultant agent models are known as Belief/Desire/Intention, or BDI, models. The beliefs, desires and intentions can then be used to generate plans of action (Bratman, Israel, & Pollack 1988).

There have been a number of attempts to define logical frameworks which capture these ideas (for instance that in (Cohen & Levesque 1990)). Perhaps the best known is that proposed by Rao and Georgeff (1991b) which models belief, desire and intention as modalities and gives them a semantics in terms of possible worlds. Whilst this is theoretically appealing, from our perspective it suffers from two drawbacks. The first of these is that there is no clear link between the model and its implementation in systems such as dMars (Ingrand, Georgeff, & Rao 1992). The second is that

despite the relation between the model and decision theory (Rao & Georgeff 1991a) the model uses a very coarse model of uncertainty (in which everything that has a non-zero probability is regarded as "possible" in the modal logic sense).

We have started to solve these two problems. We deal with the first by building agents as multi-context systems (Giunchiglia & Serafini 1994), with separate contexts for beliefs, desires, and intentions (Parsons, Sierra, & Jennings 1998). This approach allows the BDI specification of agents to be directly executed as a set of concurrent theorem provers. We deal with the second problem by associating degrees of belief (Parsons & Giorgini 1997) with individual logical formulae. At present we do this using Dempster-Shafer theory, but the argumentation framework (Fox, Krause, & Ambler 1992; Krause et al. 1995) we use to handle the degrees of belief is sufficiently general to allow for other representations of uncertainty as well, including probability and possibility theories.

This work can therefore be seen as an attempt to take a way of describing single autonomous agents, using this as a means of building such agents, and extending such agents to take advantage of existing work on decision making under uncertainty.

Social agents

Taking an agent-oriented perspective, classical decision making can be seen as a model for building agents which are utility maximisers. Such an approach allows us to build rational single agent systems, but, as Castelfranchi (1990) argues, such a view of rationality lacks a social dimension—an agent in a multi-agent community may well do better if it takes into account not only its own utility, but also the utilities of other agents in its community since (i) the agent derives some benefit from the improved performance of the community as a whole, and (ii) if a given agent helps out other agents when it is free and those other agents then return the favour, that individual agent will be able to complete its own tasks more efficiently. We have experimentally demonstrated an example of this kind of "social benefit" (Kalenka & Jennings 1997).

To capture this kind of rationality we have (Hogg & Jennings 1997) extended the classical decision making model to include the social utility of actions—the utility to other agents in the community—as well as individual utility. The validity of this model is currently being tested experimentally using a multi-agent extension of the Phoenix fire-fighting simulation (Cohen et al. 1989). This work can therefore be seen as providing a mixed-initiative extension to classical decision making in the sense that decisions are influenced by the effect they will have on other agents. However, the decisions are still made by single agents, so the mixture of initiatives is implicit in the social utility rather than being explicit in the sense that the decision is made by agents in concert.

Agents which share information

To move towards agents which engage in truly mixed-initiative decision-making, we need to be able to build agents which exchange information, and, in particular, information which leads them to reach decisions which are mutually satisfying. In agent terminology, this process is negotiation.

Our first attempt to define such agents (Parsons & Jennings 1996) involved agents making proposals for joint action and sharing information by passing arguments in support of those proposals. In this model, the shared information and the relationship between conflicting arguments allows the agents to refine their proposals until they are acceptable to all the agents involved and the community as a whole can decide to adopt them. This work was then extended by defining a minimal framework for supporting the negotiation (Sierra et al. 1997) (which is now being implemented), and by setting it in the framework of multi-context systems (Parsons, Sierra, & Jennings 1998).

As mentioned above, one of the advantages of the multi-context approach is that it makes it possible to directly execute agent models, so this work, in principle, gives us a means of building multi-agent systems which engage in mixed-initiative decision making, though the decision theory is the symbolic theory of argumentation rather than classical decision theory.

Future work

At the present time, our work on lone agents and information sharing agents relies on a non-standard theory of decision making (based on the use of argumentation) rather than on classical decision theory. One of our intentions for the future is to extend both lines of work towards the use of classical decision theory in order to build agents which are "normative" in the probabilistic sense (indeed, we have already established the conditions under which argumentation itself is normative in this sense (Parsons 1997)). However, we are wary of the naive application of decision theory, bearing in mind that the reason BDI models were initially proposed (Bratman, Israel, & Pollack 1988) was because it was felt that the resource bound on agents precluded its use. We are therefore looking at the use of qualitative approaches to decision making, including both the extension of argumentation to incorporate arguments about actions and values (Fox & Parsons 1997)

and the extension of models which use extreme probabilities for reasoning about beliefs (Bourne & Parsons 1998) towards a full decision theory along the lines of that described by Wilson (1995).

Summary

We have described three strands of our current work which together cover the spectrum of decision making agents from single agents which only consider their own desires, through agents which take account of other agents, to sets of agents which reach decisions through consultation. While our work at either end of the spectrum does not, as yet, make use of classical decision theory, we have sketched out the relationship between the classical decision theory and the model used and described how the gap between the two will be narrowed in the future.

References

Bourne, R. A., and Parsons, S. 1998. Propagating probabilities in system p. In *Proceedings of the Florida AI Research Symposium*, (to appear).

Bratman, M. E.; Israel, D. J.; and Pollack, M. E. 1988. Plans and resource-bounded practical reasoning. *Computational Intelligence* 4:349–355.

Castelfranchi, C. 1990. Social power: a point missed in multi-agent DAI and HCI. In Demazeau, Y., and Muller, J., eds., *Decentralized AI*. Amsterdam, Netherlands: Elsevier Science Publishers. 49–62.

Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. *Artificial Intelligence* 42:213–261.

Cohen, P. R.; Greenberg, M. L.; Hartand, D. M.; and Howe, A. E. 1989. Trial by fire: Undertstanding the design requirements for agents in complex environments. *AI Magazine* Fall:32–48.

Fox, J., and Parsons, S. 1997. On using arguments for reasoning about values and actions. In *Proceedings of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, 55–63.

Fox, J.; Krause, P.; and Ambler, S. 1992. Arguments, contradictions and practical reasoning. In *Proceedings of the 10th European Conference on Artificial Intelligence*, 623–627.

Giunchiglia, F., and Serafini, L. 1994. Multilanguage hierarchical logics (or: How we can do without modal logics). *Artificial Intelligence* 65:29–70.

Hogg, L. M., and Jennings, N. R. 1997. Socially rational agents. In *Proceedings of the AAAI Fall Symposium on Socially Intelligent Agents*, 61–63.

Ingrand, F. F.; Georgeff, M. P.; and Rao, A. S. 1992. An architecture for real-time reasoning and system control. *IEEE Expert* 7(6).

Jennings, N. R.; Mamdani, E. H.; Corera, J.; Laresgoiti, I.; Perriolat, F.; Skarek, P.; and Varga, L. Z. 1996. Using ARCHON to develop real-word DAI applications Part 1. *IEEE Expert* 11:64–70.

Kalenka, S., and Jennings, N. R. 1997. Socially responsible decision making by autonomous agents. In *Proceedings of the 5th International Colloquium on Cognitive Science*.

Krause, P.; Ambler, S.; Elvang-Gøransson, M.; and Fox, J. 1995. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence* 11:113–131.

Parsons, S., and Giorgini, P. 1997. On using degrees of belief in BDI agents. Technical report, Department of Electronic Engineering, Queen Mary and Westfield.

Parsons, S., and Jennings, N. R. 1996. Negotiation through argumentation—a preliminary report. In *Proceedings of the International Conference on Multi Agent Systems*, 267–274.

Parsons, S.; Sierra, C.; and Jennings, N. R. 1998. Agents which reason and negotiate by arguing. *Journal of Logic and Computation* (to appear).

Parsons, S. 1997. Normative argumentation and qualitative probability. In *Proceedings of the International Conference on Formal and Applied Practical Reasoning*, 466–480.

Rao, A., and Georgeff, M. 1991a. Deliberation and its role in the formation of intentions. In *Proceedings of the 7th Conference on Uncertainty in Artficial Intelligence*, 300–307.

Rao, A. S., and Georgeff, M. P. 1991b. Modeling Rational Agents within a BDI-Architecture. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, 473–484.

Sierra, C.; Jennings, N. R.; Noriega, P.; and Parsons, S. 1997. A framework for argumentation-based negotiation. In *Proceedings of the 4th International Workshop on Agent Theories, Architectures and Languages*.

Wilson, N. 1995. An order of magnitude calculus. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, 548–555. San Francisco, CA: Morgan Kaufman.