UBU: Utility-Based Uncertainty Handling in Synthetic Soccer

Magnus Boman, Helena Åberg, Åsa Åhman, Jens Andreasen, Mats Danielson, Carl-Gustaf Jansson, Johan Kummeneje, Harko Verhagen, Johan Walter

DECIDE/K2LAB (`www.dsv.su.se/DECIDE`)
Department of Computer & Systems Sciences
Stockholm University & The Royal Institute of Technology
Electrum 230
SE-164 40 KISTA, SWEDEN
mab@dsv.su.se

The UBU RoboCup team is described. Intuitive ideas and general objects for the participating researchers and their students are presented. The UBU team participates in the simulation league. The key idea is to repeatedly use the advice provided by a normative pronouncer (or decision module) when choosing what to do next. The pronouncer acts on input from each individual player, the basis of which is stored in a local information base. The team is under continuous development: At the time of writing, no version of UBU makes extensive use of pronouncer calls.

Team Description Format

We begin by describing the methodology of the UBU project in Section 2. The pivotal concept of pronouncer is briefly presented in Section 3, and principles of normative artificial decision making are discussed in Section 4. We then present a very coarse program architecture, together with some implementation details, in Section 5. The final section on future research is important, since we are currently in the final stages of completing the basic functionalities of the team.

Project Methodology

The DECIDE research group has since its inception focussed on normative decision analysis, and on tools for evaluation in particular (EDB96, EDB97, DE98). In recent years, some of the attention has been given to *artificial decision making* (BE95, B97, B98). As the term indicates, not only the decision makers, but the entire procedure of reaching a decision is artificial. The concept of autonomous artificial agents deciding what to do without human intervention is currently studied in several on-going

M. Asada and H. Kitano (Eds.): RoboCup-98, LNAI 1604, pp. 352-357, 1999.

projects with participants from DECIDE. These include agent assistance in securing energy contracts on a de-regulated market, agent-based intelligent building control, machine learning driven pollution control, and decision making agents in telecommunications. The projects are academic, but all have industrial participants. Some of the projects use customised simulation testbeds. The RoboCup domain is a good complement to these testbeds, since it offers a relatively simple dynamic real-time environment. Hence, we have chosen to test some our hypotheses formulated within the mentioned projects on a RoboCup team.

The UBU research team is basically a group of students and their supervisors. Magnus Boman (team captain), Mats Danielson, Carl-Gustaf Jansson, and Harko Verhagen constitute the latter category. Johan Kummeneje has recently become a graduate student in the DECIDE group with the dissertation topic directed towards RoboCup.

Pronouncers

A requirement that must be met for artificial agents to behave intelligently is that they can ask for advice. The base case is when the agent asks itself what it should do next. Nearly all AI research, as well as more than half of all agent research, can be placed in this simpler category. Many of the classical AI problems, such as the frame problem and the knowledge representation problem appear immediately, and must be addressed. The even more difficult case is when the idle agent asks someone (or something) else. This case can in turn be analysed by considering two sub-cases. Firstly, the agent may ask other agents, belonging to the same multi-agent system (MAS). Second, the agent may ask an entity that is not part of the MAS and that may in fact not be an agent at all. This entity may come in different guises, usually a blackboard, a control panel, a pronouncer, a decision module, an ontology, a knowledge repository, a daemon, or an oracle. Each of the guises just mentioned have too many variations to allow for them to be studied in precise terms: a blackboard, for instance, does not entail the same agent architecture or model to all researchers that claim to use them.

The availability of data runs from full to zero. In the former case, the entity is in some sense omniscient. Put simply, if each agent represents all its known or believed information in a knowledge base, the entity has access to a database containing the union of all such knowledge bases, with each entry typed to the agent in whose knowledge base the entry appeared. In the latter case, one must first define zero data availability. In the strictest possible sense, it means the entity accepts no input, since each input consists of data. Hence, it is solipsistic. Recalling that its sole purpose was to give advice, it is also useless. In a slightly less strict variant, zero data availability means that input may be given to the entity, but that this input reveals no information about any of the agents in the MAS. A syntactic realisation could be that input cannot contain typed or modal formulae, i.e. the entity knows neither the identity of the agent with which it communicates, nor the identity of any of the agents that might be mentioned in the input message it receives. In this case, the entity may be used as a random procedure; i.e. it can give advice of a quality equal to casting the dice.

The quality of data basically runs from precise and certain to imprecise and uncertain. The reason for saying ‘basically’ is that under some circumstances, the quality of imprecise data is equal to that of precise data. For instance, when solving a system of linear equations representing constraints on agent behaviour in order to find out what an agent should do next, imprecise data generally speaking yields a solution set of a cardinality bounded only by the number of variables in the system, while precise data yields a unique solution. A well-studied example is situations in games giving rise to multiple equilibria: it is the assumptions about the quality of data available to the players that determine whether the game has a unique equilibrium point or not.

Next, we name the entity giving advice. We study pronouncements, and therefore a candidate for naming the entity is *pronouncer*. This term is new to the computer science literature. It implies that the advice given is formal and authoritative, giving the entity a normative status, so it should be used with care, but would be appropriate for the purpose of the UBU team. Another candidate is *decider*. It is neutral, being simply the nominative form of ‘decide’. It is, however, common to used decider as short for ‘decision procedure’ in recursion theory. Moreover, it is the name of at least one commercial product in the area of risk analysis. A third possibility is *decision module*. The word ‘module’ suggests something internalised, i.e. that we are studying one module among others, intrinsic to an intelligent agent. This is less appropriate for how advice is provided in the case of UBU. The term has been used extensively in the MAS literature before, usually in connection with planning, but also for bases of heuristics in expert systems, and for software providing normative advice to an inquirer. By contrast, *oracle* suggests something extrinsic. Oracle too has been used in the MAS literature. Unfortunately, oracle already has a well-defined meaning in complexity theory. It also implies high quality of the advice given, if not omniscience. In view of the above, we will choose the decision module term.

There are two possibilities for situating the entity. One is to define a decision module as local to each agent. Just as each agent might have its own list of goals, the decision module is treated as a customised tool for decision support. Hence, the entity is not merely copied into each agent, but is adapted to the agent to which it belongs from the outset, and increasingly so during its life-span. The alternative is to have a global entity that querying agents call upon repeatedly. The entity is then a resource to be shared among the agents, and will then amount to a function, the input of which will have to carry all information about the decision situation, and the output of which will be a recommended action. This function would be centralised in the same way as a facilitator in a federated architecture (GK94). We choose the second possibility.

If our sole concern was individually rational agents, and we also relied only on the principle of maximising the expected utility, the input could be a decision tree, possibly weighted with probabilities and utilities. The function would then amount to a calculator recommending (one of) the action(s) with the highest expected value. However, we attempt to produce socially rational agents, and must therefore add group constraints, or use similar means to qualifying individually rational behaviour to achieve social intelligence (B98). This cannot be achieved by merely modifying the weights in the decision tree. Instead, such constraints are part of a local information base, with respect to which each evaluation is carried out by the decision module. The necessity of such local bases was previously realised in the context of risk constraints (EBL98): Not all risk attitudes can be modelled using decision trees.

Naturally, one can imagine simple MAS in which each agent has the same responsibility towards the group, and even in such systems non-trivial problems arise (KJ98), but the realistic and interesting case is where each agent has unique obligations towards the other agents. For instance, a MAS might consist of 200 agents in which a particular agent has obligations towards the entire population (including itself), but also towards two overlapping strict subsets of, say, 20 and 25 agents that constitute coalitions. These coalitions might be dynamically construed, something which will affect the nature of obligations heavily over time.

Interestingly enough, procedures for updating the local information base can be viewed as learning procedures. In particular, the adaptation to particular coalitions, i.e. to group constraints, can be viewed as learning how to function socially. These issues cannot be pursued in this brief team description, but are currently under investigation (BV98).

Artificial Decision Making

We make the following two provisos, more concise motivations for which are available in (B97) and (E96), respectively.

Proviso 1: Agents act in accordance with advice obtained from their individual decision module, with which they can communicate.

Proviso 2: The pronouncer contains algorithms for efficiently evaluating supersoft decision data concerning probability, utility, credibility, and reliability.

Every change of preference (or belief revision, or assessment adjustment) of the agent is thought of as adequately represented in the local information base. This gives us freedom from analysing the entire spectrum of reasoning capabilities that an agent might have, and its importance to the use of the decision module. The communication requirement presents only a lower bound on the level of sophistication of agent reasoning, by stating that the agent must be able to present its decision situation to the pronouncer, and that the agent can represent this information in the form of an ordinary decision tree, extended by general risk constraints (EBL98).

The second proviso also requires some explanations. Supersoft decision theory is a variant of classical decision theory in which assessments are represented by vague and imprecise statements, such as “The outcome o is quite probable” and “The outcome o is most undesirable" (M95). Supersoft agents need not know the true state of affairs, but can describe their uncertainty by a set of probability distributions. In such decisions with risk, the agent typically wants a formal evaluation to result in a presentation of the action (in some sense) optimal with respect to its assessments, together with measures indicating whether the optimal action is much better than any other action, using a distance measure. The basic requirement for normative use of such measures is that (at least) probability and utility assessments have been made, and that these can be evaluated.

In the local information base are non-linear systems of equations representing supersoft data about

- probabilities of the occurrence of different consequences of actions
- utilities of outcomes of different consequences of actions
- credibilities of the reports of other agents on different assessments
- reliabilities of other agents on different agent ability assessments

The preferences of the agents can be stated as intervals of quantitative measures or by partial orderings. Credibility values are used for weighting the importance of relevant assessments made by other agents in the MAS. Reliability values mirror the reliability of another agent as it in turn assesses the credibility of a third agent (E96). All bases except the utility base are normalised. Note that a MAS without norms is treated is this paper as a social structure where group utility is irrelevant to the individual agent. The presence of norms can manifest itself in various ways (B98), but unfortunately we cannot discuss this matter further here.

Notes on the Implementation

The goalkeeper of the UBU team is implemented in Java. The code was written in accordance with several key concepts of Java, e.g., threads, encapsulation, and communication. We have designed the goalkeeper agent to consist of three subsystems for communication, memory, and deliberation (or reasoning). The latter is the only role-specific code module of the goalkeeper agent. The goalkeeper is a mixed-behaviour agent in that it is reactive, e.g., in its use of the `catch` command (applied whenever the ball is within the `catchable_area`), while more deliberative in its positioning.

The field players are written in C with a separate module for basic interaction with the server including navigation, and a communications module letting each of the agents provide hints about the positions of other players and the ball. On top of these two modules, local information bases are being implemented as a 'magnets' that attract agents to areas of strategic importance. 'Negative magnetism' is used to reject agents from danger-zones (such as other players). The positions of the magnets are dynamically re-calculated, using the estimated positions of all players.

The decision module to be used is basically DELTA (D97), with certain modifications made to it. DELTA was written in C.

Future Research

The UBU team is unfinished. In Paris, UBU won one game and lost three. That version of the team had no pronouncer calls, and not even all the basic functionality was implemented. At PRICAI, a much-improved version will compete. This version will have almost full basic functionality, e.g., a reasonable treatment of offside strategies. Hopefully, it will also contain pronouncer calls. In any case, the long-term goal is the Stockholm competition in 1999: playing home has some great advantages! For that competition, a version of UBU depending highly on pronouncer calls is planned.

References

(Åb98) H. Åberg: "Agent Roles in RoboCup Teams", M. Sc. Thesis, DSV, 1998.

(Åh98) Å. Åhman: "Decision Control in RoboCup Teams", M. Sc. Thesis, DSV, 1998.

(BE95) M. Boman & L. Ekenberg: "Decision Making Agents with Relatively Unbounded Rationality", *Proc DIMAS'95*: I/28-I/35, AGH, Krakow, 1995.

(B97) M. Boman: "Norms as Constraints on Real-Time Autonomous Agent Action", Boman & Van de Velde (eds.) *Multi-Agent Rationality (Proc MAAMAW'97)*: 36-44, LNAI 1237, Springer-Verlag, 1997.

(B98) M. Boman: "Norms in Artificial Decision Making", to appear in *Artificial Intelligence and Law*.

(BV98) M. Boman & H. Verhagen: "Social Intelligence as Norm Adaptation", Dautenhahn & Edmonds (eds.) *Proc WS on Socially Situated Intelligence*, SAB'98, Zuerich, 1998.

(D97) M. Danielson: Computational Decision Analysis, Ph.D. thesis, DSV, 1997.

(DE98) M. Danielson & L. Ekenberg: "A Framework for Analysing Decisions Under Risk", *European Journal of Operations Research* **104**: 474-484, 1998.

(E96) L. Ekenberg: "Modelling Decentralised Decision Making", Tokoro (ed.), *Proc ICMAS'96*: 64-71, AAAI Press, 1996.

(EDB96) L. Ekenberg, M. Danielson & M. Boman: "From Local Assessments to Global Rationality", *Intl Journal of Intelligent Cooperative Information Systems* **5**(2&3): 315-331, 1996.

(EDB97) L. Ekenberg, M. Danielson & M. Boman: "Imposing Security Constraints on Agent-Based Decision Support", *Decision Support Systems* **20**(1): 3-15, 1997.

(EBL98) L. Ekenberg, M. Boman, and J. Linnerooth-Bayer: "General Risk Constraints", Presented at the DAS98 WS at the Intl Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria. Submitted.

(EK97) H. Engström & J. Kummeneje: "DR ABBility: Agent Technology and Process Control", M. Sc. Thesis, DSV, 1997.

(GK94) M. R. Genesereth & S. P. Ketchpel: "Software Agents", *Communications of the ACM* **37**(7): 48-53, 1994.

(KJ98) S. Kalenka & N. R. Jennings: "Socially Responsible Decision Making by Autonomous Agents", *Proc Fifth Intl Colloq on Cognitive Science*, San Sebastian, 1998.

(M95) P-E. Malmnäs: "Methods of Evaluations in Supersoft Decision Theory", unpublished manuscript, 1995. Available on URL: `http://www.dsv.su.se/~mab/DECIDE`.