UBU: Pronouncers in RoboCup Teams

Magnus Boman, Jens Andreasen, Mats Danielson, Carl-Gustaf Jansson,
Johan Kummeneje, Johan Sikström, Harko Verhagen, Håkan Younes

DECIDE/K2LAB
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. The UBU team participates in the simulation league, and is under constant development. The key idea is to use, in certain situations, the advice provided by a normative real-time decision support algorithm for letting players choose what to do next. The algorithm acts on input from each individual player, the basis of which is stored in a local information base. This idea is only implemented in part, and the version of UBU competing in Singapore is to be seen as a first step towards a team of artificial decision making soccer players.

1. Team Description Format

This document is a sequel to (Å98). We begin by describing the methodology and members of the UBU project, in Section 2. The pivotal concept of pronouncer, see (BH98), is briefly presented in Section3. Principles of 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 presents plans for future research.

2. 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).[1] In recent years, some of the attention has been given to *artificial decision making* (BE95, B97, B98, BH98). 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 decision what to do without human intervention is currently studied in several

1 More information on the DECIDE project is available on `www.dsv.su.se/DECIDE`.

on-going 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 of 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. The student group currently numbers six, but the co-authors of this report are those that have directly affected the current version of the UBU team. Jens Andreasen has implemented a full team with basic functionality in C. Johan Kummeneje, a graduate student concentrating on RoboCup, who also completed his master's thesis about agent control (EK97), has contributed to, among other things, the goalkeeper. This has resulted in a team with a goalkeeper implemented in Java, and the rest of the players implemented in C; in itself a neat example of coping with heterogeneity. Johan Sikström is currently working on details in the basic features of the team, and also preparing for a second abstraction layer in the code. This level will allow for calls to algorithms for decision support. Finally, Håkan Younes is currently surveying the usefulness of commercial as well as academic tools for artificial decision making. His aim is to provide a clear view of what is available for use by UBU, and also to do some actual test runs with various pronouncers.

3. Pronouncers

Let us first name the entity giving advice *pronouncer*. This is a new and invented term, suggesting an extrinsic entity, and also that the advice given is formal and authoritative, giving the entity a normative status. Thus, it should be used with care, but fits UBU perfectly.

There are two possibilities for situating the entity. One is to define a decision module, local to the agent. Just as each agent might have its own list of goals, such a 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 lifespan. The alternative is to have a pronouncer that querying agents call upon repeatedly. The entity is then a resource to be shared among the agents. It will 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 pronouncer would be centralised in much the same way as a facilitator in the federated architecture (GK94).

We choose this latter option, in spite of the complexity of the input to the pronouncer. If our sole concern was individually rational agents, and we also relied only on a basic choice rule generator like the principle of maximising the expected utility (PMEU), the input could be a decision tree, weighted with probabilities and utilities. The pronouncer would then amount to a calculator recommending (one of) the action(s) with the highest expected value. However, we are ultimately dealing with socially intelligent agents in UBU, 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 (EBL98). Instead, such constraints are part of a local information base, with respect to which each evaluation is carried out by the pronouncer. The necessity of such local bases was previously realised in the context of risk constraints (EBL97): Not all risk attitudes can be modelled using decision trees.

If we were to vary the evaluation rules in the pronouncer itself, e.g., to experiment with using different extensions of PMEU, it would make sense to have customised entities for normative advice, i.e. to use decision modules instead. Such pluralism with respect to decision support is easy to give arguments for in the case of individually rational agents, but it is perhaps less natural to think that individual utility maximisers are to adhere to the same norms even though their rules for evaluation differ. In any given domain, it is easier to fix a pronouncer and then vary the individual beliefs, preferences, and relevant norms, represented in local information bases.

Naturally, one can imagine simple MASs in which each agent has the same responsibility towards a group. Even in such systems non-trivial problems arise, and there it would suffice to store norms globally, as part of the pronouncer. The realistic and most general case, however, is where each agent has unique obligations towards each and every one of the other agents. For instance, a MAS might consist of 200 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.

4. 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 decision module contains algorithms for efficiently evaluating supersoft decision data concerning probability, utility, credibility, and reliability.

The first proviso makes our presentation clearer, because 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 decision module, 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 actions 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 in 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.

5. Notes on the Implementation

As mentioned at the end of Section 2, 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 used is basically DELTA (D97), with certain modifications made to it. DELTA was written in C.

6. Future Research

The UBU team competed in Paris this summer, where it won one game and lost three. The basic functionalities in that version of UBU were not satisfactory. Team play was too defensive, and UBU could not handle off-side very well. However, the basic communication platform was validated, and much was learned about necessary refinements of the basic skills. No pronouncer calls were made in the Paris version of UBU. We intend to implement at least some pronouncer calls in the Singapore version, aiming ultimately at a serious use of pronouncers at the Stockholm competition in July, 1999.

Hypotheses to be validated through more programming and testing before the Singapore contest include that the pronouncer is indeed efficient enough for real-time use, and that the updating of the local information bases is possible in the ways anticipated.

References

(Å98) H. Åberg, Å. Åhman, J. Andreasen, M. Boman, M. Danielson, C.-G. Jansson, J. Kummeneje, H. Verhagen & J. Walter: "UBU: Utility-Based Uncertainty Handling in Synthetic Soccer", *Proc RoboCup'98*, 1998.

(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*.

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

(BH98) M. Boman & H. Verhagen: "Social Intelligence as Norm Adaptation", Dautenhahn & Edmonds (eds.), *Proc SAB'98 WS on Socially Situated Intelligence*: 17-25, Centre for Policy Modelling, 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.

(EBL97) L. Ekenberg, M. Boman & J. Linnerooth-Bayer: “Catastrophic Risk Evaluation”, IIASA Report No. IR-97-045. Intl Institute for Applied Systems Analysis, Laxenburg, Austria, 1997.

(EBL98) L. Ekenberg, M. Boman, and J. Linnerooth-Bayer: “General Risk Constraints”, Presented at the DAS WS at the 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/DECIDE`.