Electronic Transactions on Artificial Intelligence
Vol. 2 (1998): *105-158*
http://www.ep.liu.se/ej/etai/1998/004/

Refereed Article

David Poole:

Decision Theory, the Situation Calculus and Conditional Plans

Author's affiliation: University of British Columbia, Vancouver, Canada.

Abstract: *This paper shows how to combine decision theory and logical representations of actions in a manner that seems natural for both. In particular, we assume an axiomatization of the domain in terms of situation calculus, using what is essentially Reiter's solution to the frame problem, in terms of the completion of the axioms defining the state change. Uncertainty is handled in terms of the independent choice logic, which allows for independent choices and a logic program that gives the consequences of the choices. As part of the consequences are a specification of the utility of (final) states, and how (possibly noisy) sensors depend on the state. The robot adopts conditional plans, similar to the GOLOG programming language. Within this logic, we can define the expected utility of a conditional plan, based on the axiomatization of the actions, the sensors and the utility. Sensors can be noisy and actions can be stochastic. The planning problem is to find the plan with the highest expected utility. This representation is related to recent structured representations for partially observable Markov decision processes (POMDPs); here we use stochastic situation calculus rules to specify the state transition function and the reward/value function. Finally we show that with stochastic frame axioms, action representations in probabilistic STRIPS are exponentially larger than using the representation proposed here.*

Publication and review history

1. First version *published* by Linköping University Electronic Press on 15.6.1998 and permanently available at
   http://www.ep.liu.se/ea/cis/1998/008/

2. *Received* by the research area "Reasoning about Actions and Change" of the Electronic Transactions on Artificial Intelligence, posted on its web site, and advertised in its Newsletter on 14.6.1998; the summary on 25.9.1998.

   Also *received* by the research area "Decision and Reasoning under Uncertainty" and advertised in its Newsletter on 3.8.1998.

3. Publicly *discussed* in those two ETAI areas during August, 1998 - June, 1999. The discussion protocol is available and will be retained in the Article Interaction Page at
   http://www.etaij.org/ra/rac/008/

4. Revised version, with minor corrections compared to the first version, *published* by Linköping University Electronic Press on 17.3.1999 and permanently available at
   http://www.ep.liu.se/ea/cis/1998/008/
   under the label "Revised publication 1999-03-17".

5. Revised version *accepted* after due refereeing and according to scientific journal standards by the Electronic Transactions on Artificial Intelligence (ETAI) on 12.6.1999.

6. Final version, with minor corrections according to the suggestions of the referees, *published* by Linköping University Electronic Press on 15.7.1999 and permanently available at
   http://www.ep.liu.se/ea/cis/1998/008/
   under the label "2nd revised publication 1999-07-15".

7. Final version appears in Electronic Transactions on Artificial Intellligence, Volume 2 (1998), Issue 1-2, pages 105–158. The issue and the annual volume are made permanently available at
   http://www.ep.liu.se/ej/etai/1998/

8. Paper editions of the ETAI issue and the ETAI volume containing this artice *republished* by the Royal Swedish Academy of Sciences.

The review policy and the quality requirements for acceptance are documented at http://www.etaij.org/info/

Article maintenance

The Electronic Transactions on Artificial Intelligence maintains an Article Interaction Page for this article at
http://www.etaij.org/ra/rac/008/

Copyright conditions

Summary

ETAI authors are recommended that each article be accompanied by a summary. Longer and more specific than a conventional abstract, it should specify in concrete terms what are the new results in the article. If present, the summary also plays a role in the refereeing process: referees are asked to judge whether the results as specified in the summary *are of importance to the field, and whether the body of the article* gives sufficient evidence *for the claims made in the summary. – The Editor.*

This paper shows how to use the situation calculus within a logic based on Bayesian decision theory. Within this logic we can axiomatize the dynamics of a domain (including stochastic frame axioms, and stochastic ramification axioms), how noisy sensors work, and the utility of final states. This paper defines the expected utility of conditional plans (containing sequential composition and conditioning on sensor values) given a representation of the domain.

The first part of this paper argues that many of the problems with representations of discrete actions have arisen because the primitive mechanisms in first order logic to handle uncertainty (namely disjunction and existential quantification) are inadequate to handle the choices (and their consequences) that an agent needs to make. For example, logics with no uncertainty, such as acyclic logic programs, seem to handle frame and ramification axioms unproblematically. We build on these acyclic logic programs, but handle uncertainty with more sophisticated tools. The paper argues that the appropriate representation for uncertainty for an artificial agent is Bayesian decision theory. In Bayesian decision theory all uncertainty is handled by probability and all preference by utility.

The representation that we use to combine logic and Bayesian decision theory is the Independent Choice Logic (ICL). The ICL grew (a) from work on abduction and default reasoning where we want to abduce to causes and use assumptions to make predictions from these, (b) from building rule-structured conditional probability tables in Bayesian networks, and (c) from work on different agents choosing assumptions modelled on game theory. The ICL is built from a choice space (a set of mutually exclusive alternatives, where an alternative is a set of ground atomic formulae, called atomic choices), and an acyclic logic program that gives the consequences of the atomic choices. The logic program can't have any clauses that imply an atomic choice. The semantics is defined in terms of possible worlds; there is possible world for every selection of one element from each alternative. A formula is true in a possible world if it it true in the (unique) stable model of logic program together with the atomic choices selected. The ICL is naturally abductive: the explanations of g form a concise specification of the possible worlds in which g is true.

In this paper all choices are made by nature. There is a probability distribution over the alternatives (a non-negative number associated with each atomic choice so the values for the atomic choices in an alternative sum to one). When there are finitely many alternatives, the probability of a world is the product of the probabilities of the atomic choices selected by the world (when there are infinitely many alternatives we use a measure over sets of worlds in the standard manner, but where the alternatives are unconditionally independent). The probability of a formula is the sum of the probabilities of the worlds in which it is true. There is a close relationship between the ground (variable-free) ICL and Bayesian networks; the elements

of the body of a rule correspond to the parents in the Bayesian network. The ICL gives rule-structured conditional probability tables.

The paper shows how the situation calculus can be embedded in the ICL, where each possible world specifies a complete history; it specifies what is true in each situation. The actions that make up the situation are the actions of the agents (that is, the motor controls sent by the agent). Thus the agent can know what situation it is in, even if it doesn't know what is true in any situation. The ICL specifies a probability distribution over fluents in each situation.

We also allow for (noisy) sensors where the sensor value is some stochastic function of what is true in a situation. The agent can condition its actions on its sensor readings. The agent then adopts a conditional plan, where the conditions are sensor reading. A conditional plan selects a final situation in each possible world, where the situation selected depends on the sensor readings in that possible world.

Utility is axiomatized as any other fluent; we assume there is a unique utility for each situation in each possible world (the utility that is obtained if the agent finishes in that situation). This can be used to define an expected utility for each conditional plan. The expected utility of a conditional plan can be computed by abduction; we explain the final utility and the sensor readings in the plan. Each explanation has an associated probability and implies a final situation, and the utility of that situation. The expected utility can be obtained directly from the explanations by computing the expectation over the utilities.

The aim of this work is not to see how many different formalisms we can put together, but to see how few we can get away with. We use a very simple logic of acyclic logic programs, a simple way to handle uncertainty in terms of independent choices, and the standard situation calculus. This lets us define the expected utility of conditional plans. It is an open question as to how far we can build on this without compromising the simplicity of the formal foundations.