ELECTRONIC NEWSLETTER ON
REASONING ABOUT ACTIONS AND CHANGE

Issue 98084 Editor: Erik Sandewall 27.11.1998

Today

Discussion don't always die easily - Hector Geffner has now come back with a reply to David Poole's answer to his original question about David's article. Well, this Newsletter is always open to the preferences of its readers.

Also, Yves Lesperance's discussion starter a few issues ago about the format of robot competitions has now received the first comment, by Mikhail Soutchanski.

ETAI Publications

Discussion about received articles

Decision Theory, the Situation Calculus and Conditional Plans

by **David Poole**

Comment from: **Hector Geffner**

David,

thanks for your answers; they helped me a lot. Just a brief follow up on the status of your proposed framework: a new model? a new language? a new algorithm? all of them? ...

Best regards. -hector

> 1. In page 24, you say "The representation in this paper can be seen as a representation for POMDPs".
>
> If so, I wonder, why not present the representation in that way from the very beginning? Wouldn't that make things much clearer?

One of the problems with writing a paper that is trying to bridge different areas is to try to explain it for all readers (or to keep them all equally unhappy). I suppose I have succeeded when POMDP researchers say "this is just a representation for POMDPs" and when situation calculus researchers claim "this is just the situation calculus with some some probability added" or those studying Bayesian networks claim that "this is just a rule-based representation for dynamic Bayesian networks with actions and observables". Describing it explicitly in terms of POMDPs may have made it easier for POMDP researchers, but not necessarily for everyone else.

I think I don't agree with this. I believe that POMDPs are a general and natural *model* for sequential decision problems that involve sensing. They are not for POMDPs researchers only; in the same sense, that logic is not only for logicians. And they are simple too (unlike some of the POMDP *algorithms* that are indeed complex).

I think that anyone dealing with *sequential decision problems that involve sensing* should know about POMDPs, whether or not they appeal to POMDP algorithms for solving them, and whether or not they deal with probabilities. I know this sounds dogmatic ... but it is the truth!!! :-)

No, really, POMDPs are very useful for understanding and identifying the different dimensions of a decision problem: transitions, costs, information; and at the same time they have little to do with probabilities (namely, even if all probabilities become zero or one, POMDPs are still very meaningful and don't reduce to anything else as 0-1 probabilities that would reduce for instance to logic).

On a more general note, I think there are three essential aspects to the work in planning and control in AI:

- models
- languages
- algorithms

Models are about the mathematics of the task: what is a problem? what is a solution? what is an optimal solution? Languages are for describing these models in a convenient way. Algorithms are for computing the solutions.

I think these three ingredients are always present in approaches to planning and control in AI, and I believe it is useful to make to them explicit; even when the algorithms may take advantage of the particular language in which the model is represented (as in Strips planning).

I understand that you expect your language to be useful not only for specifying POMDPs in a convenient way, but also for solving them conveniently. That would be great. Yet even in that case, people using different POMDP algorithms could in principle benefit from your language for setting up their POMDP models. This could also be important, and indeed, in the recent AAAI Symposium on POMDPs, the need for good languages for building POMDPs for particular applications and for exchanging benchmarks was emphasized. May be your language, as well as other action languages suitably extended, could fill up that need. As you know, we have also been doing work in that direction.

Best regards.

- hector

Debates

Competitions

From: Mikhail Soutchanski

Recently (Newsletter ENRAC 16.11, 98080), Yves Lesperance posted a summary of the discussion on how to design a robot competition that exploits robots' cognitive skills. I would like to vote for a competition (suggested by Sebastian Thrun and Martha Pollack) that is centered around the

> *theme of an intelligent office assistant. This robot would be in charge of delivering mail (which arrives sporadically), coffee (at scheduled times), and occasionally people would request tours or guidance through the building. Such a scenario emphasizes two important things: (1) on-line planning and decision making on the "cognitive" level, and (2) the integration of low-level navigation skills with high-level planning and decision making.*

My interpretation of this competition theme is the following. During the competition an intelligent office assistant has to function successfully in a dynamic, unpredictable environment. For example, if robot is delayed in a corridor, it has to increase its speed to arrive in time; if robot is going to deliver coffee to an office and finds that the door is closed, then the robot may deliver the coffee to another person who requested it; if a person took his mail, then robot has to abandon the plan to deliver that mail, etc. One of the robot's responsibilities would be to trade-off a currently performed task for another more urgent and important activity, e.g., serving coffee in time is more important than delivering mail. To give the competition the flavor of an unending task, contestants have to be prepared to serve coffee to an arbitrary number of people and deliver an unspecified number of packages (in the real competition the robots will be interrupted after some reasonably large amount of time); contestants will be given a map in advance, but the set of requests with randomly generated deadlines will be assigned at competition time. One of the attractive features of this competition is that it separates plan management issues from issues of low-level sensing (such as tasks of visual recognition or collision avoidance): the "cognitive" level may receive requests by e-mail, and recipients can acknowledge delivery by pressing a button on robot. Moreover, the "cognitive" level will need only a limited input from robot's body: 1) the current time, 2) an estimate of a current location (and, possibly, information whether a path to a destination is blocked or not).

Yves Lesperance wrote (Newsletter ENRAC 16.11):

> *Developing a test suite of tasks with varying degrees of difficulty would seem to be a good idea.*

Below, I'm sketching an additional competition scenario that some researchers may find interesting: this scenario is intended to emphasize on-line plan management and decision making.

In an office environment, there are several rooms with colored blocks inside them (e.g., all blocks are placed on white paper on the floor). The initial locations of blocks in different rooms are not given in advance: robot has to wander around to determine which blocks are located where. The robot's task is to arrange a mosaic of blocks in a designated white place; the specification of goal color patterns would be randomly generated at competition time. For example, the goal is to put in a row 3 red blocks followed by 4 green blocks or 2 yellow blocks followed by 5 green blocks (note disjunction). Assume that robots can deliver no more than 2 blocks simultaneously. Each robot participating in a competition will be either assisted or disrupted during its work by exactly the same sequence of helping or upsetting actions of a judge, but those actions will not be known to contestants in advance. The judge may take an arbitrary block from an already constructed pattern and/or put a block of a different color on a free place in the pattern. The robot has to change its plan (which blocks to deliver from which room) accordingly. The environment where the robot works also may change dynamically: doors may become (partially) closed/opened, different obstacles (possibly with colored blocks on top of them), may appear in corridors, etc. The robot that will build a goal mosaic before the deadline and faster than the other robots will win.

(The robot may either rely on vision to determine the color of blocks and the current state of a mosaic that it is building, or the robot can be told what the current state is. In the case of robots which do not have manipulators, a human assistant may put blocks on the robot which will deliver them to a goal place.)

Mikhail Soutchanski

This newsletter is issued whenever there is new news, and is sent by automatic E-mail and without charge to a list of subscribers, in plaintext format. The present postscript issue is a bona fide rendering of the same information in alternative typography. Instruction for contributors and other additional information is found at the following URL:

`http://www.ida.liu.se/ext/etai/actions/njl/`