ELECTRONIC NEWSLETTER ON
REASONING ABOUT ACTIONS AND CHANGE

Issue 98080 Editor: Erik Sandewall 16.11.1998

Today

The workshop on Cognitive Robotics at the AAAI Fall Symposium in October contained several contributions and discussions that are of great interest for the Actions and Change research area. The present Newsletter contains a position statement by Yves Lesperance where he summarizes and comments about one of the discussion topics at the workshop, namely, the format of robot competitions at the AAAI conferences.

Debates

Competitions

From: Yves Lesperance

Robot Competitions for Cognitive Robotics

It has been our experience that the winners of the AAAI and RoboCup robot competitions are those that can engineer the solution to the specific problem and who can program for long hours. **We are soliciting ideas on how to design a robot competition** that will exploit robot brains as well as robot and programmer brawn.

This was recently discussed at a panel on robot competitions held during the AAAI Fall 1998 Symposium on Cognitive Robotics. Appended below is a summary of that discussion. It is organized as follows:

1. Panel Session Summary

   3 main proposals for competition tasks made:

   - scavenger hunt task with clues (Schultz/McIlraith/Shapiro),
   - multi-skilled office assistant robot (Thrun/Pollack),
   - find out where you are and then deliver something (Baral).

2. Related Discussion in Symposium Wrap-Up

We hope that this newsletter can be a forum to continue this discussion. Please post your ideas and reactions!

Panel session summary

The 1st panelist was Alan Schultz, who is chair of the AAAI robot competition organizing committee. He gave an overview of the history of AAAI robot competitions. At first, the robots had no manipulators and the tasks assigned (e.g. navigation problems, give a tour, find life on Mars, find the remote control, clean the room) reflected this. Later manipulators became available and the tasks changed to reflect this. Vision systems are still not required.

Then, Alan said that the competition was being split into a more basic student competition, and a challenge competition where tasks require significant advances to the state of the art and stay unchanged for several years. As an e.g. of a challenge task, he mentioned "registering for AAAI, attending talks, summarizing what was learned, and putting the results in a web page". A kind of "scavenger hunt" task was also mentioned. He said the committee was very interested in entertaining proposals from the cognitive robotics community.

During the subsequent discussion, Sheila McIlraith suggested that a "scavenger hunt" task where clues are provided would be a good test for cognitive robotics. Many participants reacted enthusiastically to this proposal. Yves Lesperance pointed out that we should be careful about game-like tasks (e.g. chess), which seem to involve a specific kind of intelligence not typical of commonsense reasoning. Later in the discussion, Stuart Shapiro suggested that the task could involve finding people given clue messages in natural language. Someone also suggested that the clues be given on the web.

The 2nd panelist, Murray Shanahan, listed some criteria for a good competition task. It should not be a task where the robot becomes a mere front end to a reasoning module. It should require reasoning as part of the perception action loop. Elaboration tolerance should also be required. Erik Sandewall added to this that the task could involve performing activities while moving; it should combine planning and obstacle avoidance.

The third panelist, Michael Jenkin, said that we want to design a competition which favors 'neat' solutions over 'scruffy' ones. He suggested looking at existing 'scruffy' competitions:

- AAAI robot competition: tasks are finite, with specific and well-defined goals; teams typically hack all night and it might be said that the group with the best depth in terms of last-minute programming wins;
- RoboCup robot soccer competition: it has many categories, some involving only simulation, and some for which Nomadics and RWI have built specific robots; one concern here is that teams are quite secretive concerning the way their systems work.

Michael then suggested that we:

- use a task which is unending and requires the robot to maintain some ongoing relationship with the environment, e.g. more like ongoing mail servicing rather than 'deliver these letters'; the latter rewards solutions which do mobility well, but the former requires long term scheduling and task management in addition to mobility;
- provide a simulation-only competition class (a purely 'neat' competition);
- set up other competition classes that use standard hardware platforms; if people use, say a SuperScout, then it might even be possible for everyone to use the same hardware (which might be provided by the manufacturer).

Ian Horswill argued against doing simulation-only, claiming that in this case we wouldn't be taken seriously by the robotics community. Developing a realistic simulator was said to be as hard as using real hardware, e.g. how does one simulate robot slippage? Some questioned whether one needs to simulate such "malfunctions". Someone then made the point that participants in the competition should be forced to disclose how they solved the problem, so that the field as a whole could benefit.

The fourth panelist, Sebastian Thrun, said that existing AAAI competitions do not provide much of a challenge on the cognitive level, the reason being that they typically focus on a single, pre-determined task that does not require on-line symbolic planning and decision making. He suggested a competition 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. He encouraged people in the cognitive robotics community to really do it, i.e. build cognitive robots and compete. In the robotics community, this is the only way to get attention.

Martha Pollack then described a similar idea for a task called "Be My Administrative Assistant". Ahead of the competition, examples of possible office tasks to be assigned would be provided. Then at competition time, each robot would be assigned a randomly generated set of tasks with randomly generated deadlines, some of which may involve forming new plans out of the actions in the announced tasks; e.g.:

- make plane reservations (from the web),
- travel to the "travel agent" to purchase them,
- make copies,
- deliver them to certain people at certain times,
- answer the phone,
- etc.

Later, Chitta Baral proposed that we use some sort of office navigation task where the robot is given a topological map of the environment but is not told what its initial position is; it would have to find out where it is before working on the task. In addition, the robot may find that a certain corridor is blocked. The robot would need to recognize that, add this observation to its knowledge and make a new plan (possibly through execution monitoring) to get to its goal.

In the end, many participants seemed enthusiastic about the more ambitious tasks proposed, but some seemed to want to stick to simpler tasks where there were clear interactions between perception, action, and reasoning. Developing a test suite of tasks with varying degrees of difficulty would seem to be a good idea. But we still need to agree on a competition task.

Related discussion in symposium wrap-up – How to proceed?

Gerhard Lakemeyer raised the important point that unless we had teams that were ready to compete this summer, we shouldn't push our agenda for the 1999 robot competition.

Stuart Shapiro suggested that we perhaps start it as a (can't remember the exact term used) "demonstration sport" in the first year (like ballroom dancing at the Olympics).

It was agreed that we would try to:

- decide on a proposed new format for the RobCup, which we could announce in 1999; this would give teams time to build their systems to address the new challenges presented by the competition;
- at AAAI 2000, it would be implemented as a "demonstration sport";
- assuming success at AAAI 2000, it would be a full-fledged competition for AAAI 2001.

Note that before we proceed, we should get assurances from several researchers that they are committed to participating in this new competition.

This summary was reconstructed from my notes and material obtained from Chitta Baral, Michael Jenkin, Sheila McIlraith, and Sebastian Thrun. Apologies to anyone whose ideas may have been distorted or left out. Please post your clarifications!

Yves Lesperance
lesperan@cs.yorku.ca

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/`