

Reasoning about Actions: A Position Statement

Murray Shanahan

Queen Mary and Westfield College, London, England

Since this is a position statement, I suppose it's legitimate to give a sales pitch for a particular formalism, and to be a bit controversial.

AI needs an action formalism that is expressive, and that incorporates a solution to the frame problem that's robust in the face of the phenomena it can represent. The formalism should be expressive enough to represent at least the following phenomena.

1. Actions with indirect effects (ramifications)
2. Concurrent action
3. Non-deterministic action
4. Narrative time
5. Continuous change

A rigorous argument that the formalism in question solves the frame problem should be supplied.

Here comes the controversial bit. I believe that the most instructive way to validate a formalism is through a judiciously chosen set of representative benchmark scenarios. I sometimes feel that attempts to do this validation by proving a relationship between a formalism and some other formal structure are just an excuse for introducing a lot of unnecessary mathematics. I strongly believe that we can only contribute to AI if we divert our efforts away from proving "interesting" theorems and into the application of our formalisms to fundamental problems in AI (such as planning in robots).

Here comes the sales pitch. The event calculus presented in my forthcoming book is capable of representing all the above mentioned phenomena, and the solution to the frame problem that accompanies it is provably immune to the Hanks-McDermott problem. In recent ECAI and AAI papers, this formalism is applied to a serious example involving a robot. The example includes all five of the above listed representational features. If your favourite action formalism

can't do this robot example by now, maybe it's time to change to a new one.

Honesty time now. The event calculus cannot, as yet, handle knowledge producing actions, or complex actions (ie: actions including program constructs). To me these requirements seemed less urgent than the ones in my list, but they're undoubtedly important.

Finally, the last thing we need is new formalisms. The only excuse for introducing a new formalism is that it is fundamentally different from any of the existing ones. Instead, we have to try to expose the underlying similarities and differences between possible formalisms in the hope that we can start to understand the range of representational choices in a principled way.

The above opinions are those of the author on 20th November 1996, and may be subjected to total revision when he is older and wiser, or possibly sooner.