ELECTRONIC NEWSLETTER ON
REASONING ABOUT ACTIONS AND CHANGE

Issue 98017 Editor: Erik Sandewall 12.2.1998

Debates

Ontologies for actions and change

From: Pat Hayes

Responses to Judea Pearl and Erik Sandewall in ENRAC 26.1 (issue 98009):

Responses to Judea (all quotations from his message 26.1):

On Actions vs Observations, or on Pat Hayes' reply to Geffner, Poole and me:

> *... the cleavage between the culture that Hector, David and I represent and the one represented by Pat has gotten so DEEP that we are not even sure we are talking about the same thing.*

Indeed. I started to write a blow-by-blow response, but Judea and I trying to out-flame each other is reminiscent of a Godzilla movie, so I will ignore all the rhetoric and just try to summarise what seem to me to be the essential points of difference. First however, to clear up one basic misunderstanding:

> *Pat does not think the "distinction between sensing and acting is either necessary or even ultimately coherent". For him, observing a surprising fact evokes the same chain of reasoning as establishing that fact by external act. In both cases, so claims Pat, the world is changing, because a world is none other but one's beliefs about the world, and these do change indeed in both cases.*

No, no. Of course the world cannot be identified with one's beliefs about it. This distinction is central to my view of knowledge representation. The world is changing because that is what worlds do, by and large; certainly my own 'common-sense' world keeps changing all the time. If there were no change, our ordinary notions of time and tense would be incoherent. Again, this seems so obvious that I am sure nobody can disagree with it, so Judea and I must be using the word 'change' differently, and indeed he seems to have in mind something like a change of mind about the causal structure of the world. I'd call this a change of belief rather than a change in the world:

> *... there is no need to invoke such dramatic phrases as "changing worlds" for things that we know how to handle by standard inference methods. (e.g., adding a proposition "light on" to the system and let a classical theorem prover draw the consequences. [and if our current beliefs contain "light off" then the contradiction can be handled by either temporal precedence or minimal-change belief-revision]).*

As an aside here, I am far less sanguine about the ability of 'standard' inference methods (whatever they are) to routinely draw consequences about a causal world, and the general usefulness of such obviously false 'principles' as temporal preference or change-minimisation. (Judea seems to think that all the interesting representational work is already done. I don't know what makes him so optimistic.) But in any case, changes in the world have nothing to do with inference methods. Oddly enough, Judea here seems to be suffering from the same mistake he accuses me of, by confusing the world with our models or descriptions of it.

> *[we need to decide]... when a change is considered a "world-change" and when it is merely a "belief change".*

Let me see, how can I possibly say this more clearly? A "world change" is a change in the world, and a "belief change" is a change in belief. My point is precisely that the time which the beliefs are *about* need bear no simple relation to the time at which the beliefs are *created.* (This is similar to the distinction in temporal databases between 'valid time' and 'transaction time'.)

> *Question: Is flicking the switch and seeing the light come on an example of a changing world ?*

Yes.

> *From a physical viewpoint, the world has obeyed Schrödinger's equation before and after the action, so it has not changed at all.*

This remark seemed to come out of left field until I re-read Judea's messages several more times, and then I suddenly got it. Judea's entire framework is the Copenhagen interpretation of quantum mechanics. This is where the otherwise mysterious talk of Schrödinger, the 'laws of astrophysics', invariants, the need to establish 'boundaries', the sharp distinction between actions and observations (which, in the Copenhagen interpretation, must come from 'outside' the system boundaries, causing a 'collapse' of the wave-function), the Niels-Bohrian talk of "all the particles in the universe", and the odd idea that its observations which change the 'world', all come from. It also may account for his doctrinal confidence that his rather idiosyncratic perspective is accepted by most philosophers (anyone who thinks that **anything** is accepted by most philosophers obviously has never been in close association with philosophers.) OK, if you think of our subject as likely to be informed by quantum theory, then go ahead. Personally I disagree, for various reasons which I wont take space up here enlarging on.

One last comment:

> *... we are not dealing here with physics or psychology. We are dealing with various formal systems for modeling the world and our beliefs about the world.*

To deny both physics and psychology, and then to immediately speak of "world" and "beliefs about the world", seems to me to be at best careless.

Responses to Erik

In his message 26.1, referring to an earlier message by myself, he wrote:

> In the sitcalc (any variety), actions are changes in the world, not motor commands. One plans by thinking about the changes, not by thinking about the muscles one is going to use. Putting such a system into a robot requires one to somehow connect these actions with motor controls, no doubt, but they shouldnt be identified. (Murray and Ray, do y'all agree??)

However, earlier in the same contribution Pat had written:

> ... Our peripheral systems often blend motor action and lowlevel perception in tight feedback control loops, so that our bodies seem to 'move by themselves', but these lowlevel controls are the result of more cognitive decision-making (deciding to hit a tennis ball, say.)

Pat, I can't make sense out of your position: at one point you seem to argue that low-level and high-level descriptions of actions can't ever be separated; at another point you seem to say that they are best treated in complete separation.

These two passages were in different contexts. I dont think that observations and actions can be clearly separated, ultimately, basically because observation processes often get so integrated into our actions that the combination is best described together in terms of feedback and control, as when we talk of 'moving carefully', or 'holding rigidly in place'. I bet such mergings of observation and action will occur in robotic applications pretty soon, in fact. But this is quite consistent with my other point, since the action/observation combinations can be described and reasoned about at a higher level than muscle signals, just as simpler actions can; and this higher level - ie in terms of their relations to the world rather than to anatomy - is usually the most appropriate for action reasoning. In fact, in much of higher-level planning, they may be treated similarly as 'actions', except that some actions require some observational capability. For example, consider the fact that moving something carefully depends on being able to see where it is, and inferring that careful movement might require switching the light on (but a coarse job can be done in the dark.)

My own preference is to take both into account, but to be precise about having two distinct levels of descriptions with distinct roles. In particular, this allows for dealing both *with an action as a prescription for motor controls* and *as the expectation for what state changes will be obtained as a result. It also allows one to relate those two levels to each other.*

Well, I agree. I guess my original response (to David Poole) was a reaction to his assumption that the 'muscle' case was the normal one. But maybe I over-reacted. I'm quite happy with making a distinction like Erik's careful one between "material" and "deliberative" levels, or something like that, but he has misunderstood me slightly:

> *... The deliberative level is the one that Pat alludes to, where actions are characterized by discrete properties at a small number of time-points: possibly only the beginning and the end of the action, possibly a few more, possibly a sequence of partial world states at integer time-points (as in basic Features and fluents).*

I'm beginning to despair of being understood. Didnt this entire newsletter discussion start with me saying that I **don't** want to endorse the discrete-time-point way of describing change? Ive written extensively on time-intervals, and the ontology I used for naive physics was based on extended pieces of space-time ('histories') rather than anything pointlike.

The fact that someone with Erik's insight seems to be unable to think of 'high-level' as meaning anything other than instantaneous situations is a vivid illustration of the problem Ive been complaining about: almost all our thinking in this field is dominated by the archaic and awkward oversimplifications embodied in the situation calculus.

David Poole wrote:

> What if I didn't know whether there was something on x when I tried to pick it up? It seems that the only sensible interpretation of the precondition is that if there was nothing on x and I carried out $pickup(x)$, then the expected thing would happen. If something was on x and I carried out $pickup(x)$ then who knows what may happen. The role of the precondition is that it is only sensible to attempt to carry out the action when the preconditions hold.

Pat Hayes answered:

> No no. This is all a *description* of an action. .....

To which Erik Sandewall answered:

> *I'll agree with David if the last sentence is changed to go "...it is only sensible to attempt to carry out the action when the preconditions* are believed to *hold". .....*
>
> *Therefore, Pat, I don't understand what you are objecting against in this respect.*

I have no objection to your rephrasing, but I dont think its a trivial change: that's exactly what I was reacting to. The distinction between a fact in the world and a belief in the planner is an important one when we are having a discussion about action and observation.

> *(And what does it have to do with the original need for nonmonotonic logic?)*

Because when an observation contradicts a belief, we would seem to have contradictory beliefs, unless our logic is nonmonotonic. This was the original motivation for McCarthy's use of circumscription, to allow the reasoner to use implications of the form

$$\text{if } Property(state) \text{ then } Otherproperty(do(action, state))$$

even when there could clearly be circumstances where, in fact, if the action were taken, the conclusion might not be true (the block might not be held if the grippers were oily). If we know ahead of time that actions always come out the way we want them to, then we could write change axioms which were guaranteed to be correct, and circumscription wouldnt be necessary (just as it often isnt necessary when describing programs, where it is possible to give a *complete* description of the relevant parts of the 'world'-state - where 'world' here means the computer the program is planned to run on, of course.)

Im puzzled that you need to ask this question. Isnt this textbook-level stuff?

Pat Hayes

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/`