ELECTRONIC NEWSLETTER ON REASONING ABOUT ACTIONS AND CHANGE

Issue 98034 | Editor: Erik Sandewall | 13.4.1998

Today

The discussion about ontology of time continues with an additional contribution by Pat Hayes. This discussion has so far been included under the general discussion on the topic of "ontologies for actions and change", but now it seems more appropriate to spin it off into a topic of its own. This is motivated both by the number of significant contributions that have been made, and by the general importance of this topic.

Thus, for readers of the E-mailed newsletter, everything proceeds as before, but the structured discussion record in our webpage structure has been revised retroactively to correspond to the new classification.

The present issue also contains a question by Rob Miller to Erik Sandewall re his ETAI submitted article.

ETAI Publications

Discussion about received articles

Additional contributions have been received for the discussions about the following article(s). For the full context, please refer to the on-line structure which contains one article interaction page with the entire discussion for each article.

Logic-Based Modelling of Goal-Directed Behavior

by Erik Sandewall

Comment from: Rob Miller

Dear Erik,

I'm confused about axioms S4, S5 and S8 involving the action $test(p)$ on page 7 of your paper:

S4 seems to state that it is always possible to perform a test to see if a given property is true. But I would have thought that the opposite is nearly always the case - that in most commonsense or robot environments we could almost assume by default that the agent is not in a position to directly test the truth of a given property of the environment. (This is certainly true for me as I type this question at home. I can't check the condition of anything in my office, in the other rooms in my house, etc. without first performing some other actions.) So what is the rationale behind axiom S4? Is the agent perhaps "testing" it's own knowledge, rather than the real environment?

S5 states that test actions are always instantaneous. Why should this be so, at least in a framework where other actions are not instantaneous?

S8 says that, once executed, a test action fails if and only if the corresponding property is false. But this seems to go against the commonsense notion of "testing". For example, when I test to see if I have any new emails (actually by clicking on the "?" button on Netscape) I count this action as successful if the system reports "no new messages on server". I count the action as failed if (as frequently happens with my faulty modem here at home), clicking on "?" crashes my computer. I'm then left with no information as to whether I have new emails or not. So is this not a test action in the sense that you intended?

Regards,

Rob

Comment from: **Erik Sandewall**

To:**Rob Miller**

Dear Rob,

The "action" *test*(p) is classified as an action in the formal sense that it can be combined with other actions using, for example, the sequential operator. In those cases that you describe, where the actual performance of the test takes a nontrivial amount of time compared to other events, I imagine that one would use *test*(p) merely for the internal decision within the robot and using knowledge that it has already acquired. The corresponding, external action of finding out what is the case and *then* succeeding or failing would then be represented with something along the lines of *findout*(p) ; *test*(p), where *findout*(p) is non-instantaneous.

Will *test*(p) always be immediately preceded by *findout*(p) in practice? I think of at least two kinds of situations where this will not be the case. When some fluent is being monitored continuously by a lower software level, then the current value of p can safely be assumed to always be available. Also, in those cases where the agent has to plan ahead and make sure *in advance* that it acquires important information, so that it is there without delay when needed, then again *test*(p) will not be *directly* preceded by and tied to *findout*(p).

The axioms assume that the argument of *test*(p) is either true or false. The case where incomplete knowledge is allowed may be handled by any of the well-known methods, for example by having a pair of two propositional fluents p and p', and where lack of information is represented by making both of them false.

The mnemonic "test" was inspired by an opcode in an assembly language that I learned a long time ago.

Regards, Erik

Debates

Ontologies for time

From: Pat Hayes

Answer to Jixin's contribution to this discussion on 1.4:

There is really little point in arguing about which theories are more 'intuitive' unless one is more precise about what one's intuitions are. There are two fundamental problems with arguments like this. First, intuitions are malleable, and one can get used to various ways of thinking about time (and no doubt many other topics) so that they seem 'intuitive'. Second, our untutored intuition seems to be quite able to work with different pictures of time which are in fact incompatible with one another. Jixin's own intuition, for example, seems to agree with McCarthy's that time is continuous, and yet also finds the idea of contiguous atomic 'moments' (intervals with no interior points) quite acceptable. But you can't have it both ways: if moments can meet each other, then there might not be a point where the speed is exactly 60mph, or the ball is exactly at the top of the trajectory. If time itself is discrete, then the idea of continuous change is meaningless. Appealing to a kind of raw intuition to decide what axioms 'feel' right lands one in contradictions. (That was one motivation for the axiom in Allen's and my theory, which Jixin found "awkward", that moments could not meet. The other was wanting to be able to treat moments as being pointlike. That was a mistake, I'll happily concede.)

Jixin wrote:

> *What follows is our response to the arguments about the ontology of time from Pat Hayes, Ray Reiter, and John McCarthy.*
>
> *Response To John:*
>
> *The example of car accelerating demonstrates the need of time points for time ontology.*
>
> *A similar example is throwing a ball up into the air. The motion of the ball can be modelled by a quantity space of three elements: going-up, stationary, and going-down. Intuitively, there are intervals for going up and going down. However, there is no interval, no matter how small, over which the ball is neither going up nor going down. The property of being stationary is naturally associated with a point, rather than any interval (including Allen and Hayes' moment), a "landmark" point which separates two other intervals.*

Yes, I agree. However, notice that there is a coherent frame of mind which would deny this. According to this intuition, which is similar to Newton's old idea of the infinitesimal, one would say that there are no points, but some intervals are so small that they can be treated like points at a sufficiently larger scale. In this perspective, it would be false to claim that there was *no* interval at which the velocity is zero; rather, one would say that the interval was infinitesimal. (If you want to deny the reasonableness of this perspective, first reflect on the fact that it is nearer to physical reality than any model based on the real line.)

> *Response To Ray: (—)*

I agree with Jixin here.

> *Response To Pat:*
>
>> ...For example, the 1990 AIJ critique of Allen's account by Galton (wrongly) assumes that Allen's intervals are sets of points on the real line.
>
> *After re-reading Galton's paper* [j-aij-42-159], *as we understand, Galton's arguments are in general based on the assumption that Allen's intervals are primitive, rather than sets of points on the real line. In*

> *fact, the main revision Galton proposes to Allen's theory is a diversification of the temporal ontology to include both intervals and points. That is, in Galton's revised theory, intervals are still taken as primitive.*

Galton's intuitions are clearly based on thinking of intervals as sets of points. He takes it as simply obvious, for example, that there is a distinction between open and closed intervals.

> *Having pointed out this, however, as shown in Ma, Knight and Petrides' 1994 paper* [j-cj-37-114]*, Galton's determination to define points in terms of the "meeting places" of intervals does not, as he claims, axiomatise points on the same footing as intervals, and hence that some problems still remain in these revisions.*
>
> > there does seem to be a simple, basic, account which can be extended in various ways to produce all the other alternatives, and this core theory is the one I was referring to.
>
> *Does this core theory refer that one in which "intervals are uniquely defined by their endpoints (which are also the points they fit between) and two intervals meet just when the endpoint of the first is the startpoint of the second"? Or Allen's one? - It seems the former one.*
>
> *Anyway, yes. There* **does** *seem to be such a simple, basic core theory. For general treatments, in Ma and Knight's CJ94 paper* [j-cj-37-114]*, a time theory is proposed (as an extention to Allen and Hayes' interval-based one) which takes both intervals and points as primitive on the same footing - neither intervals have to be constructed out of points, nor points have to be created as the places where intervals meet each other, or as some limiting construction of intervals. The temporal order is simply characterised in terms of a single relation "Meets" between intervals/points.*

This theory seems to be similar to that outlined in my 1990 paper with Allen, [j-ci-5-225] (and given at greater length in a U of Rochester tech report of the same date.) But there is little point in bickering about who said what first, as almost all this discussion (including for example Allens 13 relations) can be found in publications written in the last century, if one looks hard enough. All AI work in this area (including my own) is like children playing in a sandbox. The theories and idea themselves are a much more interesting topic.

One technical point, about 'primitive'. One of the things I realised when working with James on this stuff was that if ones axioms about points were minimally adequate it was trivial to define interval in terms of points; and one can also define points in terms of intervals, although that construction is less obvious. (I was immensely pleased with it until being told that it was well-known in algebra, and first described by A. N. Whitehead around 1910.) Moreover, these definitions are mutually transparent, in the sense that if one starts with points, defines intervals, then redefines points, one gets an isomorphic model; and vice versa. So to argue about which of points or intervals are 'primitive' seems rather pointless. We need them both in our ontology. If one likes conceptual sparseness, one can make either one rest on the other as a foundation; or one can declare that they are both 'primitive'. It makes no real difference to anything.

> *Some advantages of this time theory are: (1) It retains Allen's appealing characteristics of treating intervals as primitive which overcomes the Dividing Instant Problem.*

See above. But in any case this doesnt overcome the problem. Allen's treatment allows lights to just come on, but it doesnt provide anywhere for the ball to be motionless.

> *(2) It includes time points into the temporal ontology and therefore makes it possible to express some instantaneous phenomenon, and adequate and convenient for reasoning correctly about continuous change. (3) It is so basic that it can be specified in various ways to subsume others. For instance, one may simply take the set of points as empty to get Allen's interval time theory, or specify each interval, say T, as ⟨T-left, T-right⟩ where T-left $\leq$ T-right, to get that one Pat prefers.*

Not quite right. In my simple theory, *T-left* isnt before *T-right*, it equals it. ....

> *Yes, it's true.* **And**, *it seems that, all these can be reached equivalently by simply taking pointlike interval $\langle t, t\rangle$ as identical with point t in the case where both intervals and points are included in the time ontology.*

Yes, that is another alternative way to formalise things.

> one can, for example, say something like
>
> ```
> ((illuminated i) or (dark i)))  implies
>     (not (pointlike i))
> ```
>
> so that the light is neither on nor off **at** the switching point. In this theory, every proposition has a 'reference interval' during which it is true, and a proposition might not be true of subintervals of its reference interval. (Though some propositions might be. This kind of distinction has often been made in the linguistic literature. Note however that this intuition is basically incompatible with the idea that an interval is identical to the set of the points it contains.)

> *This can be distinguished by applying Holds_In and Holds_On (that is Allen's Holds, see Galton's 1990 paper* [j-aij-42-159]*).*

Yes, exactly, although there is no need to use this formal strategy, as I explain in the time catalog section 1. Briefly, HoldsIn P i is true just when i is a subinterval of a reference interval j where HoldsOn P j . Again, it is largely an aesthetic judgement, but I find Galton's *hold-on* vs. *hold-in* distinction awkward and unintuitive. (It suggests that there are two different 'ways to be true'.)

> *....Actually, as shown in Ma and Knight's 1996 paper* [j-cj-37-114]*, to characterise the intuitive relationship between Holds-On and Holds-In, in the case where intervals are allowed, some extra axiom is needed.*

I will check this paper to see what you mean in detail, thanks.

> *Also, it seems that, in Pat's formulation, for expressing that interval $\langle a, b\rangle$ is a subinterval of interval $\langle c, d\rangle$, one would have $c \leq a \leq b \leq d$. In this case, we get that pointlike intervals $\langle a, a\rangle$ and $\langle b, b\rangle$ (or equivalently, points a and b)* **fall in** *interval $\langle c, b\rangle$.*

Yes; but note that if the theory uses reference intervals, that fact that P holds for an interval I doesnt imply that it holds for every point (still less every interval) **in** I. So this is quite consistent:

- Ball is rising for interval $\langle a, b\rangle$
- Ball is falling for $\langle b, c\rangle$
- Ball is stationary for $\langle b, b\rangle$

You can consistently add that rising and falling are true for all nonpointlike subintervals and every properly contained subinterval of the reference interval.

> *On the one hand, many cases suggest the need of allowing a proposition to holds at a single point. For instance, see the example of throwing a ball up into the air described earlier in the response to John MaCarthy.*
>
> *On the other hand, allowing a proposition to holds at a single point doesn't necessarily make pointlike un-definable. It depends on if one would impose some extra constraints, such as*
>
> ```
> ((illuminated i) or (dark i)))  implies
>       (not (pointlike i))
> ```
>
> *as introduced by Pat for the light switching example, which actually leads to the assertion that the light is neither on nor off AT the switching point.*

In my theory it leads to the conclusion that $\langle a, a\rangle$ does not exist (or, is not a reference interval for 'illumination'), where a is the switching point. All the distinctions between kinds of point - ones where something is true and ones where something is switching - can be cast into a typology of intervals. (This example illustrates why I like to distinguish between the point a - which undoubtedly exists, is where the intervals meet, has a clock time, etc. - and the interval $\langle a, a\rangle$, which, if it existed, might be an embarrassment.)

> *Actually, in the later version of Allen and Hayes's theory that appears in 1989* [j-ci-5-225]*, an awkward axiom is proposed to forbid moments to meet each other. It is interesting to note that, although moments are quite like points (moments are non-decomposable), they still have positive duration (they are not pointlike). Moments are included in Allen and Hayes' time ontology, while points are not. One of the reason that such an axiom is awkward is that it doesn't catch the intuition in common-sense using of time.*

I agree. This was awkward in our old paper, and stemmed from our reluctance to accept the idea of intervals which could meet themselves. I'm now reconciled to that idea: in fact, it seems inevitable, much as the existence of zero seems inevitable once one allows subtraction.

> But these cases only make sense if one thinks of interval and points in the usual mathematical way, which is exactly what Im suggesting we don't need to do. We can get almost everything we need just from the ordering structure: we don't need to get all tied up in distinguishing cases which can only be formally stated by using all the machinery of real analysis.

> *The cases make sense* **not** *only if one thinks of intervals and points in the usual mathematical way. In fact, all the three cases are demonstrated under the assumption that both intervals and points are treated as primitive, rather than in the usual mathematical way.*

Well, it depends on what axioms one assumes! Perhaps I have been speaking too carelessly about the 'usual mathematical way'. Heres my intuition: the standard account of the continuum seems *forced* to resolve the dividing point problem by deciding which interval contains the point, distinguishing open from closed intervals, because it *identifies* an interval with a set of points. (So if both intervals 'contain' the point, the intervals must intersect.) One can take points as basic or intervals as basic or both as primitive; that's irrelevant, but the crucial step is that (set-of-points = interval) identification. Thats exactly what I want to avoid. My point is only that if we abandon that idea (which is only needed for the formal development of analysis within set theory, a rather arcane matter for us), then there is a way to formalise time (using both intervals and points as primitive, if you like) which neatly avoids the problem.

> ... Is the light on or off at (exactly) 3.00 pm? The only way to answer this is to find a suitable non-pointlike interval of light or darkness completely surrounding 3.00 pm, because 'being on' is the kind of proposition that requires a nonpointlike reference interval.

> *But it seems that there are also some other kind of proposition to which one cannot assign any nonpointlike reference interval. For instance, in the throwing ball up into the air example, proposition "the ball is stationary" can only be true at points, and for any point we cannot find any non-pointlike interval (completely) surrounding it over which the ball is stationary.*

Yes, exactly. Some properties can accept pointlike reference intervals, some can't. Like Galton's distinction between 'at rest' and 'motionless' (former requires nonpointlike, latter doesnt.) Thats the point. Notice the distinctions are now all about intervals. They arent between different ways of being true, but are bread-and-butter distinctions between intervals, expressible within the theory. The machinery of truth wrt an interval is the same in both cases (and in others, eg 'intermittently true' and other exotic variations.)

> This has nothing to do with whether an interval is open or closed: in fact, there is no such distinction in this theory. It only arises in a much more complicated extension which includes set theory and an extensionality axiom for intervals.

> *In the time theory where both intervals and points are taken as primitive, we can (if we like) talk about the open and closed nature of intervals with some* **knowledge** *being available. This kind of knowledge can be given in terms of the Meets relation, rather than some "much more complicated extension which includes set theory and an extensionality axiom for intervals". In fact, we can define that:*
>
> - *interval I is left-open at point P iff $Meets(P, I)$*
> - *interval I is right-open at point P iff $Meets(I, P)$*

- *interval I is left-closed at point P iff there is an interval I' such that $Meets(I', I) \wedge Meets(I', P)$*
- *interval I is right-closed at point P iff there is an interval I' such that $Meets(I, I') \wedge Meets(P, I')$*

That's all, and it seems quite intuitive......

It is important to note that the above definition about the open and closed nature of intervals is given in terms of only the knowledge of the Meets relation. However, if one would like to specify intervals as point-based ones, such a definition will be in agreement with the conventional definition about the open and closed intervals.

That certainly seems to be an elegant device. (Though the definitions have nothing to do with knowledge; all Jixin is saying is that the definitions of open and closed can be given in terms of *MEETS*. As Allen and I showed in our old paper, the entire theory can be reduced to *MEETS*.) However, in order to be nontrivial, it must be that points 'separate' meetings, ie if $meets(I, P)$ and $meets(P, J)$ then $\neg meets(I, J)$, right? For if not, all left-open intervals are also left-closed, etc. This seems to make 'points' similar to our old 'moments': in fact, if Jixin's theory predicts $meets(P, Q) \rightarrow P = Q$ for points P and Q, then I'll lay odds it is isomorphic to our moments theory. One of the main observations in our paper was that with the no-meets axiom, one can map moments to points with no change to the theorems provable.

On the other hand, if the theory allows distinct points to *MEET*, I'd be interested to know how it is able to map smoothly to a conventional account of the continuum, since that is provably impossible. One-point closed intervals exist everywhere on the real line, but no two of them are adjacent. Atomic adjacent times (whatever we call them) are pretty much a definition of discrete time models, and are incompatible with density, let alone continuity.

Pat Hayes

PS. Maybe the most useful thing would be to put all these axiomatic theories into some common place with a common syntax - we could use vanilla-KIF - so people can compare and contrast them. I dont have enough, er, time to offer to do this right now, im afraid, but will cooperate with anyone who will volunteer.

References:

[j-aij-42-159] Anthony Galton. *A critical examination of Allen's theory of action and time.* Artificial Intelligence Journal, vol. 42 (1990), pp. 159-188.

[j-ci-5-225] James F. Allen and Patrick J. Hayes. *Moments and points in an interval-based temporal logic.* Computational Intelligence, vol. 5, pp. 225-238.

[j-cj-37-114] Jixin Ma and Brian Knight. *A General Temporal Theory.* Computer Journal, vol. 37 (1994), pp. 114-123.

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/`