ELECTRONIC NEWSLETTER ON
REASONING ABOUT ACTIONS AND CHANGE

Issue 98067 Editor: Erik Sandewall 11.9.1998

Today

Today, we present Ray Reiter's invited review of David Poole's contribution earlier this summer. Also, one more question by Paulo E. Santos to Murray Shanahan, answer by Shanahan to Baral, and answer by Ternovskaia to Denecker et al.

ETAI Publications

Discussion about received articles

Decision Theory, the Situation Calculus and Conditional Plans

by **David Poole**

Comment from: Ray Reiter

David,

I enjoyed reading your paper, and learned a lot from it. I agree with you that axiomatic action theories need an account of probability/decision theory, although I do not share your reservations about being able to do so in the sitcalc (see below). My comments and questions concern only the probabilistic (as opposed to the decision theoretic) components of your paper.

Technical issues

1. Examples 2.8 and 2.9.

The predicate $pickup\_succeeds(S)$ in the first clause.

(a) It is not parameterized by the thing – in this case key – being picked up. It should be $pickup\_succeeds(key, S)$. Presumably, you need one of these predicates for each action type, e.g. *putdown_succeeds*, *goto_succeeds*, etc. Therefore, why not introduce a generic predicate $succeeds(action, S)$, in which case, you can write $succeeds(pickup(key), S)$, $succeeds(goto(loc), S)$, etc? Similarly, you can have a relation $fails(action, S)$. Now, C_0 contains $union(succeeds(pickup(key), S), fails(pickup(key), S))$, where the union is over all ground instances of these atoms. Incidentally, your specifications of what sets are in C_0, e.g. bottom p. 13, is an abuse of notation, since it looks like an object level sentence, when in fact it is a metalevel statement for which universal quantification over situations is not defined.

(b) In axiomatizing at on p. 20, you do not include a *goto_succeeds* atom in the clause body. Does the atom $would\_not\_fall\_down\_stairs(S)$ play this role?

2. In the conditional plan of p. 21, don't you want $at(key, r101)$ instead of at_key?

3. Consider the axiom of Example 2.9 and your comment "Note that this implies that putting down the key always succeeds". Suppose not. Then is it the case that the axiom should be modified by replacing the atom $A \neq putdown(key)$ by something like

$$\neg(A = putdown(key) \wedge putdown\text{-}succeeds(key, S))$$

If so, perhaps this should be pointed out. (See my call below for a guide for the perplexed.)

4. In the axiom of Example 2.9, you include a choice $keeps\_carrying(key, S)$, which you introduce because "there may be a probability that the agent will drop the key". You then axiomatize this with $P0(keeps\_carrying(key, S)) = .95$ But I can't think of any realistic way to axiomatize this probability when only the parameter S is available.

Basically, you want to describe the probability that the robot doesn't lose the key during the interval between the occurrence of the last action in S, when $carrying(key, S)$ was true, and the occurrence of the next action A. In other words, you want to describe the probability that the robot is still carrying the key at situation $do(A, S)$ given it is carrying it in situation S. But that depends, at the very least, on WHEN this next action A occurs. Now on one account of adding time to the sitcalc (Reiter, Proc KR'96), an action can take an additional temporal argument denoting that action's occurrence time, for example, $openDoor(t)$. Then one expects that $P0(keeps\_carrying(key, S))$ given that $openDoor(10)$ is the next action should be greater than this probability given that $openDoor(1000)$is the next action. So as far as I can tell, the choice $keeps\_carrying(key, S)$ should take another argument, namely A, or perhaps simply $time(A)$, where $time(A)$ is the occurrence time of the next action A.

The need for action preconditions

On p. 8, you claim "None of our representations assume that actions have preconditions; all actions can be attempted at any time." I'm sceptical of this claim. Consider your example domain, where you adopt the axioms $P0(pickup\_succeeds(S)) = .88$ Just prior to this, you say "$P0(pickup\_succeeds(S))$ reflects how likely it is that the agent succeeds in carrying the key *given that it was at the same position as the key and attempted to pick it up*". So you can't seriously write $P0(pickup\_succeeds(S)) = .88$ because here S is any situation, including one in which the agent is 1000 kms from the key. I take it that what you intended to write was something like

$$at(key, Pos, S) \wedge at(robot, Pos, S) \rightarrow P0(pickup\_succeeds(S)) = .88(*)$$

If this is the case, then it is exactly here where you are using action preconditions, namely, in characterizing those situations in which an action has a high probability of succeeding. It also appears that you are appealing to action preconditions in the axiom of Example 2.8 (namely, $at(robot, Pos, S) \wedge at(key, Pos, S)$). Probably, these atoms should be omitted because their effects are realized in the above axiom (*).

Finally, you often talk about actions being "attempted", as in the above quotation. I found this particularly confusing since nothing in your ontology deals with action "attempts". Actions simply realize their effects with certain probabilities. You do not, for example, have an action *try_pickup*, as do other authors (e.g. James Allan) who do try to axiomatize action *attempts*. In this connection, I had a hard time making sense of footnote 9, since it motivates the axiom of Example 2.9 in terms of talk about "trying" to pick up the key. In fact, even after reading this material several times, I couldn't figure out footnote 9 until, much later, I began to think about how your approach compares with Bacchus, Halpern and Levesque. See my comments below on BHL.

Section 3.3

I must confess that I couldn't make much sense of the discussion in Section 3.3, especially your arguments that the sitcalc is inherently unsuitable for modeling multiple, possibly reactive agents. It's true that Golog lacks the right control structures for doing this, but your comments on p. 26 and p. 27 (p. 26 "In order for a situation calculus program to ... and does the appropriate concurrent actions." p. 27 "When there are multiple agents ... or complex actions.") suggest to me that you are not aware of more recent work by De Giacomo, Lesperance and Levesque on Congolog that addresses precisely these issues (Proc. IJCAI 97). Congolog provides reactive rules (interrupts) and interleaving concurrency in a semantically clean way, all formulated within the sitcalc. I believe that Congolog addresses all your criticisms of the sitcalc, except, of course, it does not (yet) incorporate utilities.

Comparison with Bacchus, Halpern and Levesque

So far as I know, there have only been two proposals for augmenting the sitcalc with probabilities, yours and BHL. As I indicate below, these two approaches are considerably different, and therefore I believe it is important to elaborate on these differences more than you do (top p. 25) especially the substantial ontological differences.

For BHL, there is no *action_succeeds* and *action_fails* fluents. Instead, an action whose outcome depend on nature is represented by the nondeterministic choice of simpler, deterministic actions corresponding to each different outcome. For example, the nondeterministic action *flipaCoin* is represented as the complex action

$$flipaCoin = flipHeads|flipTails$$

where | is GOLOG's operator for nondeterministically choosing between the two actions *flipHeads* and *flipTails*. For your ongoing example, $pickup(key)$ would be represented by

$$pickup(key) = pickup\_succeeds(key)|pickup\_fails(key)$$

When an agent executes the action $pickup(key)$ in situation S, nature steps in and selects which of the two actions $pickup\_succeeds(key)$ and $pickup\_fails(key)$ actually occurs, and it does so with frequencies determined by an axiom of the form:

$$P0(pickup\_succeeds(key), S) = p \leftrightarrow (suitable\ conditions\ on\ S\ and\ p)$$

In other words, your atomic choice of bottom p 13, and your axiom for $P0$ top of p 14 are represented by BHL in the way I have indicated above. These are very different ontological and representational commitments than yours, and I think they deserve a deeper analysis than your discussion provides. Here are a few important differences these commitments lead to:

1. Primitive actions for BHL are entirely different than yours. BHL primitive actions would be things like $pickup\_succeeds(key)$, $flipHeads$ etc, i.e. they consist of the deterministic actions that nature can perform, together with those deterministic actions (if any) that the agent can perform. For you, primitive actions are typically nondeterministic, and correspond to those actions that the agent can perform (but with nondeterministic outcomes chosen by nature), e.g. $pickup(key)$, $flipaCoin$.

2. Therefore, for BHL, situation terms describe the "actual" history selected by nature in response to the agent performed actions, e.g.

$$do(pickup\_fails(key), do(flipHeads, S0)).$$

For you, situations denote the history of agent performed actions, e.g.

$$do(pickup(key), do(flipaCoin, S0)).$$

The BHL axioms all refer to the former situations. In your ontology, only the latter are allowed.

3. In BHL, all primitive actions ($pickup\_succeeds(key)$, $flipHeads$) are deterministic and successor state axioms (the causal laws + solution to the frame problem) are formulated only wrt these. For you, the causal axioms are formulated wrt the (nondeterministic) agent's actions ($pickup(key)$, $flipaCoin$); hence the need for your fluents $pickup\_succeeds(S)$ and $pickup\_fails(S)$. These function as random variables that determine the effects of the nondeterministic action $pickup(key)$. O ne consequence of this difference is that for BHL, successor state axioms are perfectly straightforward, whereas for you, because of these random variables, the causal laws become, at least conceptually, somewhat opaque (e.g. footnote 9). Moreover, BHL use "classical" action precondition axioms, in contrast to you (see above about action preconditions).

4. For BHL, your nondeterministic agent actions like $pickup(key)$, are not part of the language of the sitcalc. Rather, they are nondeterministic GOLOG programs. Since you also introduce a notion of a program (Sections 2.8, 2.10), there are some natural comparisons to be made. The obvious (and perhaps only) important difference is in the nature of the primitive program actions. For BHL programs, the primitives are the deterministic actions described above (e.g. $pickup\_succeeds(key)$, $flipHeads$) whereas for you, they are the agent actions like $pickup(key)$.

Suggestion: A Guide for the Perplexed

Your paper uses a single, ongoing example to illustrate the axiomatization style required by your approach. Personally, I couldn't abstract, from the example, a set of guidelines by which I could invent my own axiomatization for a different example of my choosing. I think it would considerably enhance the paper if you provided such guidelines. For example, for each action, tell the reader what axioms need to be specified. What choice of fluents needs to be made? (For example, there seem to be at least two categories of fluents – "normal" fluents like $carrying(X, S)$, and atomic choice fluents. How does one decide on the latter, when given the primitive actions and "normal" fluents of an application domain?) Similarly, for each fluent, what axioms need to be specified? Also, what needs to be axiomatized about the initial situation? Finally, what do all these axioms look like syntactically? In connection with this last question, notice that on p. 20, the clause with head $at(X, P, S)$ resembles neither an effect axiom (as in Example 2.8) nor a frame axiom (as in Example 2.9).

An Inductive Definition Approach to Ramifications

by Marc Denecker, Daniele Theseider Dupré, and Kristof Van Belleghem

Comment from: Eugenia Ternovskaia

Dear Marc, Daniele, and Kristof,

Thank you for an interesting discussion regarding your paper. It clarified some subtle points of your work and helped to understand your motivation.

> *An interesting question generated by your remark is whether Thielschers approach to derive causal rules from state constraints (about one state) and dependency information can be generalised to "constraints between states" (plus dependency information as far as needed).*

This is probably a question to Michael. I found his idea about generating some of causal rules interesting, and I am also curious whether it is possible to generalize it.

Best regards,
Eugenia

A Logical Account of the Common Sense Informatic Situation for a Mobile Robot

by Murray Shanahan

Comment from: Murray Shanahan

Chitta,

Many thanks for your comments and questions.

> *(i) I think that in general a robot does not really sense events, rather it notices changes in certain fluent values (in say resistors linked to the sensors) from which it abduces the occurrences of triggered events. Nevertheless, it is ok with me to skip this part and directly talk about robots observing triggered events. But a clarification would help.*

But isn't a change in the voltage passing through a sensor an event? Especially if we consider that voltage passing through a threshold value.

> *(ii) Since you mention several times (abstract, Section 2, etc.) that you use a novel solution to the frame problem using 'Releases' it will be nice if you say about it a little more than what is said in Section 2 (just after EC 5). (You do use it in axioms in later sections, but you don't discuss them.)*

The use of "Releases" isn't really a big feature of the solution. More important is the division of the theory into parts which are circumscribed separately. This is what I call "forced separation" in my book, "Solving the Frame Problem" (MIT Press, 1997). There's a lot of detail about this solution to the frame problem in that book, which makes me a bit reluctant to duplicate it here, as it really is a side-issue. Maybe I should add a few sentences if things aren't clear. But perhaps the following answers will suffice.

> *Although from the illustrations I can appreciate the use of 'Releases' in formulating continuous change, its utility in formulating constraints is not clear to me. For example, why not eliminate B4 and instead of 'Releases' have 'Terminates' in the head of rule E2.*

Yes we could do that. But only because I've abstracted away the continuous motion in the Rotate action, and made it instantaneous. If the Rotate action initiated a period of continuous motion, then we would have to use Releases, as in the case of the Occupies fluent. This is because the axioms enforce the following rule:once a fluent has been initiated (terminated) directly by an event, only another event that directly terminates (initiates) it can change its value.

> *To me the advantage of having constraints such as B4 is that by having it we avoid explicit compilation of it to effect axioms. I.e. we can use constraints for automatically deducing effects indirectly. But if we have both B4 and E2 then we are not really avoiding the compilation, as E2 is like an effect axiom.*

Actually the purpose of (B4) here is to cut out spurious models. (But again, this only applies if we represent the continuous motion properly, instead of abstracting it away as I have done here.) Without (B4), during the period of continuous rotation of the robot, it could be facing in many different directions at once.

> *(iii) I am not sure if 'Bearing' is a standard geometrical notion. Perhaps an intuitive meaning of it will help. Or may you can use the more familiar notion of 'slope'.*

The word "bearing" is used in navigation. That's why I chose it. Maybe it's not the mathematically most appropriate term.

> *(iv) You say (just after Sp3): "the term Line(p1, p2) denotes the straight line whose ...". Perhaps it is more appropriate to say "straight line segment" instead of "straight line".*
>
> *(v) I think in axiom (B3) you are assuming velocity to be one. If you do please mention it.*
>
> *(vi) The sentence after (B6). In it you explain Blocked(w1,w2,r) by 'if object w1 cannot move any distance at all in direction r without*

overlapping with another object'. Perhaps you should say: "if object w1 cannot move any distance at all because of w2 in direction r without overlapping with another object'.

I'll make the changes you suggest. Thanks.

(vii) When defining HoldsAt(Touching(w1,w2,p),t) are you making some assumptions about the shape of w1 and w2. Imagine two triangles which touch at a point, you can align them such that they touch but yet do not satisfy the conditions you have in B6.

Yes, I think you're right. But the definition in (B6) is correct for the sense of touching required here, which excludes cases like the one you mention. I should probably add a sentence on this.

(viii) I am not clear about the intuition behind the notion of partial completion in Section 5; especially in the text after (5.7). For example in proposition 5.8, I would encode the intuition "bump switches are not tripped at any other time" by

$$Happens(Switch1, t) \rightarrow t = ind(T, bump)$$
$$Happens(Switch2, t) \rightarrow t = T_{\{}bump\}$$

I am not clear about the intuition behind your formulation.

I think your formulation is equivalent to mine. I just elected to separate the if and the only-if halves of the completion.

Also, why not use the standard Clark's completion and explain the Clark's completion of Ψ, where Ψ is as mentioned in Definition 5.9 and the paragraph before. I.e, have the Clark's completion of Ψ in the right hand side of the turnstyle in Proposition 5.8

I can't use the Clark completion, because I only want to complete the Happens predicate for sensor events. I might have incomplete knowledge about other events.

(ix) In section 6 you logically express a region as a list of straight lines. (You say it just before Bo3.) Is there any particular reason you use lists instead of sets. If not, since you already use set notation in the rest of your formulation, by using sets here also, you might avoid additional axioms such as Bo3 and Bo4

I don't have a good answer to this. I guess I could have used sets, and the result would have been slightly more succinct.

(x) Also (Bo3) is a fixpoint expression and you probably need (as needed when defining transitive closure) to minimize 'Members' to get the right models.

(Recall that the logic program

$$anc(X, Y) \leftarrow (X, Y)$$
$$anc(X, Y) \leftarrow (X, Z), anc(Z, Y)$$

is not equivalent to the formula

$$anc(X, Y) \leftrightarrow (X, Y) \vee ((X, Z) \wedge anc(Z, Y))$$

)
You are not minimizing 'Members' in Proposition 6.2. Or perhaps I am missing something.

There is an axiom missing, which says that nothing is a member of Nil. I'll put this in the final draft.

But given that, I don't think there's a problem here. It's easy to prove, for example, that $\neg Member(A, Cons(B, Cons(C, Nil)))$, just by recursing down the list (assuming A, B, and C are distinct). It's not the same as the transitive closure case.

(xi) Although I can appreciate the usefulness of the logical formulation in Section 2-7, some may pose the following question: Why not just formulate as in Section 8 (with some extensions perhaps) and then have the correctness of the algorithm in Section 9 proven with respect to the formulation in Section 8. Why go through the formulation in Section 2-7? I think it will be a good idea to address this or say a few lines about this to preempt such questions/attacks on logical formulation.

This is a very good question, and the issue is fundamental to our community's approach to AI. I don't think I can answer it fully here. But one argument would appeal to the fact that perception is a process which involves knowledge and reasoning, and if we're going to use logic for knowledge representation and reasoning in "higher level" cognitive processes, then we should use it throughout. I think this argument becomes more forceful when we consider less trivial sensors than the bump switches used here.

I hope that answers all your questions. Many thanks for taking so much trouble over my paper.

Murray

Comment from: **Paulo Eduoardo Santos**

Thank you very much for answering my first questions. I really think it is a very interesting paper.

The second bunch of questions follows:

1. In the last answer set we read:

By splitting the circumscription into parts, we get a more managable theory, one whose consequences are easier to work out. In fact, the role of circumscription here is mainly to complete certain predicates, specifically Initiates, Terminates and Happens. If we circumscribe everything together, we have to do a lot more work to avoid difficulties like the

I agree that Separation avoids the Hanks-McDermott problem , but doing so you get a restriction in the theory, since you are not able to write formulae using Initiates (or Terminates) and Happens together (on the contrary it won't be clear where to place such formulae in the circumscriptions). Thus, we are not able to express (and infer) some facts.

Then, why do not use negation-as-failure instead, since, for the particular case of circumscription we are interested in, the former is equivalent to the later (as proved in [Gelfond-Przymusinska-Przymusinski]) ?

In which sense circumscription is more powerful than negation-as-failure in this framework?

2. I could not understand the last formula of page 14. There a biimplication is used defining the region g with the formula Π. My interpretation of this is that Π is not an approximation of g but it describes exactly g. Is that what it means?

3-

In the beginning of page 15 it is considered, as a solution to the abduction process, "conjunctions $M2$ of formulae" such that a given sensor data (Ψ) is logical consequence (considering the background theory).

My question is: what do you do if there is no $M2$ that explains Ψ ? (I'm not considering the case where $M2$ follows from the background theory).

Many thanks for this opportunity,

Paulo Eduardo Santos


```
@Article{Gelfond-Przymusinska-Przymusinski,
  author =   "M. Gelfond and H. Przymusinska and T. Przymusinski",
  title =   "On the relationship between Circumscription and
               Negation as Failure",
  journal =   "Artificial Intelligence",
  year =   "1989",
  volume =   "38",
  pages =   "75-94"}
```
