ELECTRONIC NEWSLETTER ON REASONING ABOUT ACTIONS AND CHANGE

Issue 98057 Editor: Erik Sandewall 17.7.1998

Today

The discussion about ETAI submitted articles is the most salient activity in this Newsletter. Today we have a number of such discussion contributions:

- Murray Shanahan, continued comment about the article by Denecker, Theseider Dupré, and Van Belleghem;
- The authors of that article, answer to questions by Ternovskaia and by Thielscher
- Rob Miller, question to the article by Cervesato, Franceschet, and Montanari
- Erik Sandewall, answers to questions by three anonymous referees.

As usual, the full discussions can be recovered from the article interactions pages of the respective submissions, and the best off-line reading is obtained by making a printout of the postscript version of the present newsletter.

ETAI Publications

Discussion about received articles

Logic-Based Modelling of Goal-Directed Behavior

by **Erik Sandewall**

Comment from: **Erik Sandewall**

To:**Anonymous Referee 1**

> *Generally, the paper has the flavor of a progress report rather than a well-knit article. My understanding is that this perfectly complies with the publication policy of ETAI. However, sometimes such a status of a paper is indicated by adding the qualification "preliminary report" to the title, which the author might consider in the present case as well.*

My understanding is that the subtitle "preliminary report" is used in conference publications in order to distinguish an article from a later, revised version of the same article that is to be published elsewhere. The present version of the article is intended to be the final one, except for the corrections that follow from the review comments.

> *The author himself admits that there is still a lot of work to be done. In particular, a number of rather complex axioms are proposed, of which it is claimed that they characterize "obvious properties" (p.5) of relation symbols which have been introduced with a certain meaning in mind. Past experience taught us that it can be dangerous to propose axiomatizations which seem most natural at first glance. I therefore strongly support the author's remark on the need for an "underlying semantics, and a validation of the present approach with respect to that semantics..." (p.16).*

See the first item in the answer to Anonymous Referee 3.

> *An essential, most basic property of any useful axiomatization is consistency. Maybe the author could add a result in that direction.*

I guess this will be a natural corollary of validation results relative to an underlying semantics.

> *p.7: The purpose of Section 2.6 is not clear. Composite fluents aren't used elsewhere in the paper, and so one wonders, for instance, why logically equivalent composite fluents cannot be treated as equal.*

Composite fluents are used in the examples in section 2.5 (although it's possible that order between those two sections ought to be reversed). Treating logically equivalent composite fluents as equal does not have any obvious advantage, and it will complicate things considerably for the occlusion predicate.

> *p.8: The operators G and D_s have been introduced as abbreviations for formulas, e.g., $G(s, a) \leftrightarrow H(s, inv(a))$. Why is it, then, that they cannot serve as part of axioms, provided the axiomatization includes the respective definitional formulas?*

This is for reasons of convenience rather than logical necessity. The tentative specifications in question are easy to follow when they are written in terms of the abbreviations, but not when the abbreviations are expanded. Furthermore, verifying the important properties of the specifications becomes quite messy. All of this becomes clearer if the specifications are written without the use of the abbreviations.

> *p.8: The predicate Composite is introduced but (as far as I could see) not used elsewhere in the paper.*

It is used in axiom S2 on page 5, where it indicates one of the cases where an action does not start executing from the time it is invoked. The reason for making this exception is that an action of the form $a; b$ can not be allowed to start executing at a time when the action a is already executing. Therefore, it is not sufficient to say that an invoked action starts executing unless *the same action* is already in the midst of execution.

Comment from: **Erik Sandewall**

To:**Anonymous Referee 2**

> *Although I'm attracted to the author's work, I am tempted to ask why we should want an axiomatic account of goal-directed behaviour. For example, Axioms G1 to G4 really constitute an algorithm, and perhaps make better sense presented as an algorithm. From my point-of-view, the role of logic in cognitive robotics is to represent the world, not to represent the internal workings of the agent. Perhaps the author would like to address this question, in an electronic discussion and/or in the final version of the paper.*

My argument goes in two steps:

1. Rational agent behavior in a robot needs a precise specification, using logic
2. Such a specification can gain considerably by being integrated with a logic of actions and change.

First, rational agent behavior is fairly complex in the case of a robot that interacts with the real world. It requires pursuing goals, choosing plans, making appropriate reactions to new information, dealing with uncertainty and temporal constraints, and so on. Even though in the end this behavior is going to be realized by software, it is clear that a high-level specification of that software will be necessary for the actual software engineering as well as for making it possible to validate that software and to prove its properties.

Furthermore, if this software is described in a purely algorithmic style, it is difficult to avoid introducing a lot of administrative machinery. This is exemplified by the article "A Formal Specification of dMARS" by Mark d'Inverno et al. (Proceedings of the ATAL-97 workshop, pages 155-176), where a variant of PRS is specified in the Z notation (which is a software specification language). In my approach, the agent's rational behavior is specified in terms of restrictions: for each scenario, the set of models for the formalization is intended to equal the set of rational behaviors, each interpretation being one possible history in the world. Such a constraint-oriented approach has chances of being much more concise and to the point.

Finally, once we decide that it is appropriate to specify rational robotic behavior in logic, we also have to deal with actions. The robotic behavior consists of actions, so this already defines a connection point. Also, when several agents are involved, the actions of one need to be observed and understood by the others. In the longer perspective, I think we will see a lot more integration between theories for reasoning about actions and for characterizing agent behavior.

> *In Section 2.5, the third formula on page 7 entails that the postcondition of an action is false until the action is finished. But consider the example of a compound action WashDishes, which has the effect of making CutleryClean true and CrockeryClean true. Now, it's bound to be the case, assuming one item is washed at a time, that either CutleryClean becomes true before CrockeryClean or the other way around. Either way, one of the action's postconditions becomes true before the whole action is finished.*

The text preceding the formula was:

> As another simple example, consider the case of actions which are described in terms of a precondition, a prevail condition, and a

> postcondition, where the postcondition is at the same time the termination condition for the action...

The formula in question therefore refers to a limited class of actions, rather than to my approach as a whole. (Besides, in the case you mention, the postcondition of the compound action ought to be the conjunction of the two conditions, that is, $CutleryClean \wedge CrockeryClean$. Each action is assumed to have *a postcondition*, in the singular, in the example).

> *It's not clear to me what it means for the robot to "state" a formula on page 9. Does this mean it passes that formula to a control module?*

Yes. This deductive process, which one may think of as a control module, is modelled as receiving formulae from two sources: both from a higher software layer of the robot itself ("decisions") and from the surrounding world via the sensors ("observations"). Both kinds of "signals" may influence actions of the agent. This is discussed in section 3.5, the last paragraph on page 12.

> *I have one serious presentational objection to the paper. I think it would be very much improved if there were some concrete examples showing what inferences can be drawn from the axioms. (By a concrete example, I mean one with meaningful fluent names.) Even a single concrete example illustrating some of the main ideas would be of enormous benefit to the reader's intuition. Otherwise, a hostile reader might think this is all just empty formalism.*

Such an example is forthcoming.

> *Here are some more minor presentational suggestions. At the start of Section 2.6, the author introduces logical connectives for composing fluents. But these have already been used in Section 2.5. Perhaps this material could be reshuffled. Similarly, on page 10, the author declares how the variable symbols a and g will be used, although they have already appeared many times in the paper.*

This is easily arranged (although there were in fact good reasons for the order in which this material is presently presented).

Comment from: **Erik Sandewall**

To:**Anonymous Referee 3**

> *On the negative side, the paper does not really prove that its formalism is good for any specific purpose. It does not even make any mathematically precise claim of this kind. It would be good to include a description of goal-directed behavior in a toy domain and prove that some intuitively expected conclusions abour goal-directedness do indeed follow from the given axioms using the entailment methods proposed by the author. This would be more convincing than the "hand-waving" arguments in favor of the proposed approach given in Sec. 4.*

A proof that correct conclusions are obtained for one or two toy domains does not prove that the formalism is good for any interesting purpose either. As indicated in subsection 5.1 ("Limitations of these results"), it remains to validate the proposed approach using a plausible underlying semantics. This is in line with the methodology for research in this area that I have defended at a number of occasions, beginning in my IJCAI-93 paper. Only so can one identify the range of applicability of the approach in a reliable way.

You might say that the paper ought not to be published until such a validation exists and can be included. This would be an excellent position to take, provided that it were shared by the present community. However, even a quick look at the literature in our field will show that that's not the way things work: Citations to earlier work refer almost exclusively to the *approaches* proposed by the earlier author, and quite rarely to the *theorems that were proved* in the earlier paper.

Besides the formal validation, it's the applications that will prove what the formalism is good for, but this would also take too long in the present paper.

> *In the absence of such an example, the paper is reminiscent of the work on actions done in the early days of AI, such as the "frame default" in Reiter's 1980 paper on default logic, or the formalization of the blocks world in McCarthy's 1986 paper on circumscription. The ideas were interesting, but their authors were unable to prove anything about them. As the familiar shooting scenario demonstrated, a nonmonotonic formalization that looks plausible may turn out to be unsatisfactory after all. If the author of this paper tries to check that his theory works for one or two toy examples, he may very well discover bugs that need to be fixed.*

I have of course done a few toy examples for my own sake; the details are too long and boring to merit going into the paper. With respect to the historical comparison: the bugs in those cases were in the entailment method rather than the axiomatization. The approach described here is not dependent on the exact formulation of the axioms, and would not collapse if some of the axioms were to require debugging. Compare subsection 5.1.

Another comparison with the work on the frame problem in the 1980's is more appropriate: those were the days of first exploration of the problem (besides for a small number of articles in the 1970's), and it is quite natural and respectable that one started by looking for approaches that were at least intuitively plausible. *After that start* it was natural to look for precise and proven properties of those approaches. Ramification and causality have followed suit, and qualification is in its early stages. The characterization of rational robotic behavior is just in its very first stage of development, that is all.

> *It seems to me that Rob Miller was right when he suggested in his message to the author that test(p) is the action of testing the agent's knowledge rather the real environment, and that otherwise the axioms for testing are not convincing. The notation proposed in the paper uses the same symbol p to represent the fluent itself and the agent's knowledge about it, which looks peculiar.*

Well, it wasn't Rob who suggested it; you refer to my answer to his question. Anyway, the full characterization of a rational agent is bound to be fairly extensive, and the distinction between the actual and the believed value of a fluent is one of the aspects of the full problem. I believe it is best to address this complex problem by divide-and-conquer, so several other papers including references [11] and [12] "complement the work reported here", as stated in the second paragraph of section 1.2. Those articles use distinct predicates for actual and estimated feature values, so the problem you mention has indeed been addressed in the context of the present approach and we have shown how to deal with it.

The present paper works with the simplifying assumption that the agent has perfect knowledge of the fluents. This is in particular in the interest of the reader, since it makes for heavy reading to address all problems at once.

> *Regarding the author's suggestion that the incompleteness of knowledge be represented by distinguishing between p and p′ "where the lack of information is represented by making both of them false". I have this question: How would you represent the assertion that p is true but this fact is not known to the agent?*

By making both p and p' false; this is a standard technique. However, I don't recognize having mentioned it in this paper?

> *Re Theorem 1: There seems to be an implicit assumption here that the set of intervals in question is finite. Shouldn't it be included in the statement of the theorem?*

No, the proof does not require that assumption.

> *Re Theorem 2: I am puzzled by the use of the word "conversely" here. It seems to me that both parts say more or less the same thing.*

The theorem has the form "In all the models for the axioms, P. Conversely, Q". It is correct that Q follows from P even without the use of those axioms, and the second part of the theorem is merely for showing the conclusion from another point of view. Maybe it would be better to write it as a separate comment.

> *In the 3rd displayed formula on p. 7, conjunction is applied to fluents, which is only explained in the next section, and an interval is used as the first argument of H which, as far as I can see, is not defined at all.*

These bugs will be corrected.

The Complexity of Model Checking in Modal Event Calculi with Quantifiers

by Iliano Cervesato, Massimo Franceschet, and Angelo Montanari

Comment from: Rob Miller

Dear Iliano, Massimo and Angelo,

I've several questions about your paper "The complexity of Model Checking in Modal Event Calculi with Quantifiers" submitted to the ETAI.

Generally, I would like to understand the relationship between the modal logic event calculi you describe and refer to in your paper and the classical logic versions of the event calculus that people such as myself and Murray Shanahan use. So any remarks you have in this respect would be helpful. Here are some more specific questions related to this:

- Do you see your formalisms as what we describe as "narrative" formalisms - i.e. formalisms in which the flow of time is described independently of any actions that occur?
- Is the underlying concept of time linear? By "partial order" of events, do you mean "incomplete information about a total order", or are your techniques applicable to genuinely partially ordered time structures (such as the forward branching time structure employed by many versions of the Situation Calculus)?
- I'm confused about the syntax of your language(s). Are there any meta-predicates in your language other than "br"? Perhaps some examples would help. How roughly might you express information such as:

```
"Pushing causes the door to open if it's not locked"
"The door was pushed at 9 o'clock"
```

  (In the classical logic variants of the event calculus, you'd write formulas such as

```
(FORALL t).[Initiates(Push,Open,t)  <-  NOT Holds(Locked,t)].
Happens(Push,9).
```

Apologies if answers to these questions are already to be found in the other papers to which you refer. As a proponent of "the event calculus" I regret not knowing more about your previous work.

Rob

An Inductive Definition Approach to Ramifications

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

Comment from: Murray Shanahan

Marc, Daniele and Kristof,

Many thanks for your answer to my question about your paper. I think it's an excellent paper.

I've just completed a short paper on the same topic, which you may be interested in. I've submitted it to the AIJ. Basically it shows how the event calculus from my book "Solving the Frame Problem" can be used to tackle the ramification problem, including tricky examples like Thielscher's circuit. The paper is available via the following URL.

`http://www.dcs.qmw.ac.uk/~mps/ramifications.ps.Z`

All the best, Murray

Comment from: authors

To:**Eugenia Ternovskaia**
Dear Eugenia,

Your comments are very useful. We didn't have a close reading of your paper yet, but a first look showed that there is strong correspondence between our approaches. We have had very similar intuitions. Thanks for the careful reading and the corrections.

> *Perhaps the main question to you is how the third truth value would be represented in the two-valued setting?*

We think there is no need to represent the third truth value in a 2-valued setting. Observe that the transition function is a partial mapping from 2-valued states (and sets of actions) into 2-valued states. IF an expert designs his effect theory such that for all states which satisfy the state constraints and sets of actions which satisfy the action preconditions, the transition function is correctly defined, (i.e. the model of the rule set is 2-valued and no contradictory fluent literals are caused), and IF this effect theory is represented correctly in situation calculus, the transition of that state and those actions will be defined and will be 2-valued. So, in this case, a third truth-value is not necessary.

In your approach, you impose a syntactical restriction (no cycles) which guarantees that the effects are unique and well-defined. Our approach does not impose a syntactical restriction, but there is still the methodological constraint that the expert should design his effect theory such that it leads to well-defined transitions (at least for the above mentioned states and sets of actions). In case this methodological rule is followed, it is not necessary to introduce a third truth value in situation calculus or event calculus. This methodological rule requires from the expert that he establishes the 2-valuedness of the well-founded models of his effect theory at least in all states satisfying state constraints and actions sets satisfying action preconditions.

This is a point that deserved more attention.

> "Unless nondeterminism is explicitly introduced, our theory leaves no room for ambiguity."
>
> I do not think this is an advantage of your theory. I think the truth value "u" should be obtained in two cases.
>
> First, it may be a result of a "bad axiomatization". Second, it may appear due to lack of complete information about the current state of the world.
>
> (It seems that you apply the Closed World Assumption in your definition of a state as a set of (positive) fluent literals, if I understood correctly.)
>
> I believe your definitions might be easily changed to represent the second case as well. For example, a state could be associated with a set of "known" positive and negative literals.

You propose to use the truth value "u" for two different purposes: as a pointer to ill-defined cause literals and as a pointer to unknown atoms. This may be problematic. When a cause literal is 3-valued, how does the expert know whether to interpret it in one way or the other?

In our view, "u" is reservated for pointing to ill-defined causal literals. Uncertainty caused by nondeterministic actions must be modeled explicitly through nondeterministic effect rules and is modeled by having different transition states (see the section about this topic). We do not model uncertainty on the input state; we do not need to do so. We expect that in any approach which integrates our effect theories in a full temporal language (e.g. EC or SC), uncertainty on the state at a certain time point is represented by having different models with different states at that moment. In each single model, the state is well-defined and can be represented as a 2-valued set on which our transition function can be applied. In our experience, this is a better way to model uncertainty than through belief sets. We can't see in what way we would be applying the Closed World Assumption.

> It also seems strange that you cannot derive anything about "holds", although I understand your motivation. Would it be difficult to introduce this feature?

This is not really true. The result of the transition function defines the holds predicate in the next state. One can verify that the next state is derived from the previous state and the causal literals as described by the normal situation calculus persistence axiom.

> p. 23, footnote 17: A more or less meaningful "real-world" example could be something like this: an object cannot be at location p and q at the same time (but must be somewhere), and actions $move_{to_p}$ and $move_{to_q}$ are always applicable.

We agree that the footnote may not be correct, but it's not clear to us what you mean by that example. Another example of a theory with negative loops is given in the Murray Shanahan's variant of the circuit example. Though we consider the semantics of the rule set as ill-defined, the rules are certainly meaningful.

> *p.38-39 "This approach contains several contributions: - A complete uncoupling of ramifications from state constraints..."*
>
> *I do not think they can be uncoupled* completely. *Even in your theory, ramifications impose some restriction on your function Trans, I believe. Trans maps states and actions to states. In this sense reachable states are restricted and state constraints are implied. Can you clarify?*

This remark is related to your other remark:

> Also, in your summary you say "the standard view on causal laws as rules serving to restore the integrity of state constraints is too restrictive." From talking to different people, I never had the impression that this is the only and standard view on the role of causal rules in the community.

First, let us clarify what we mean by uncoupling: we define the semantics of the effect rule theory without taking state constraints into consideration. In many other approaches, this is not the case: in fact in most approaches causal rules directly imply a state constraint, they are a state constraint augmented with some directionality-of-propagation information; an exception to this is Michael Thielscher's approach, but there state constraints are used in the calculation of the possible successor states, in that effect propagations may stop as soon as a state satisfying the constraints is reached. In our approach the successor state calculation is in no way dependent on the constraints. It is always carried out to its end and yields a unique successor state. However, it is true that the end result of the calculation is checked against the state constraints. These constraints may cause rejection of the found successor state, but will in that case not calculate another (valid) successor state.

Observe moreover that ramifications need not imply state constraints even if they restrict Trans. Of course a restriction of Trans rules out certain transitions, but that does not need to mean that any states are ruled out: a state is ruled out only if all transitions leading to it are invalid. The "counter" example in our paper gives an example where this is not the case: whenever an output changes to true, the counter must be augmented. But a resulting state in which the counter were not updated (but all other effects applied) need not be invalid in itself. This state might well have existed at an earlier time in the history of the network. It can only not be reached from the currently given state. So ramifications need not imply any state constraints.

It is absolutely true that in many cases, there is a relation between state constraints and ramifications; in that case the method of Michael Thielscher is applicable and works fine in our approach as well (we applied it in the circuit example; see our answer to Murray Shanahans question). In this case, the set of effect rules should satisfy the following correctness criterion: that any state satisfying the state constraints and any set of actions satisfying the action preconditions and occurring in this state lead to a transition state in which the state constraints are satisfied. This is not automatically satisfied but should be proven in our approach.

Also we discussed with a number of people about the issue of the relationship between state constraints and effect rules; had anybody contradicted us on this point, we would not have put it the way we did. But we have no problem in weakening the point. Still, it seems true that in most approaches, state constraints are underlying or are directly involved in describing the semantics of causal rules. Maybe this is not a deliberate, conscious choice, but in the existing formalisations it is apparently a fact.

Marc, Daniele, Kristof