Electronic News Journal on Reasoning about Actions and Change 58
Vol. 2: *58-58*, 28.2.1998
http://www.ida.liu.se/ext/etai/actions/nj/9802/

Protocols from Workshop Discussions:

Workshop on Logics of Commonsense

Queen Mary and Westfield College, January 7-9, 1998

Edited by: Erik Sandewall
Linköping University, Sweden

The **Fourth Symposium on Logical Formalizations of Commonsense Reasoning**, or **Commonsense 98** for short, took place from the 7th to the 9th of January, 1998; it was organized by Rob Miller and Murray Shanahan at Queen Mary and Westfield College in London, England. The workshop program consisted mostly of presentations of accepted research papers, in the usual format, but the workshop participants decided to also try a novel way of extending the workshop discussions using the ENRAC.

The following procedure was used. During the workshop, the present editor took notes from the discussion after each presentation of a paper. These notes were written up and circulated to the participants during the rest of January, so as to offer an opportunity for correction and extension. On January 31 they were posted on-line and offered for additional commentary for the readers of the ENRAC Newsletter. (Workshop participants who were not already ENRAC subscribers were offered to join).

Two features were obtained in this fashion:

1. An opportunity for authors to add more substance to their answers than they may have done during the question period at the workshop;
2. An opportunity for all workshop participants to continue the dialogue with the authors, and with each other around the topics defined by the workshop papers.

The following pages contain the discussion protocol for the period from the actual workshop through the end of February. Those contributions where no date has been specified, are summaries of what was said at the workshop – necessarily brief in most cases. When a date is indicated, on the other hand, it indicates a written contribution by the workshop participant in question.

The full papers can be obtained from ENRAC's webpage structure, which contains links to the current state of the discussion and to the author's home pages as well as links to the actual articles.

Electronic News Journal on Reasoning about Actions and Change
Vol. 2: *59-62*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/001/

Protocol of the discussion at the FCS-98 workshop and of
the subsequent on-line discussion during February, 1998
about the following research article:

Grigoris Antoniou, D. Billington, and M.J. Mahler

Sceptical logic programming based default reasoning - Defeasible logic rehabilitated

Presented at the workshop by Grigoris Antoniou

Q1. **Bob Kowalski**

How is your representation related to Pollack's theory of argumentation?

A1. **Grigoris Antoniou** (2.2)

We have not yet compared DL to Pollock's work, but comparison to argumentation frameworks is next on our agenda.

Q2. **Michael Thielscher**

The conclusions that your logic gives for example 4.5 (a modified Nixon diamond example) seem to indicate that your logic is plain wrong. The conclusions that you obtain are not in line with common sense.

A2. **Grigoris Antoniou** (2.2)

The statement that the outcome is "plain wrong" is questionable. This debate has already been held in the context of inheritance systems. Where "a clash of intuitions" exists nobody can claim that one particular approach is THE common sense approach.

What may cause some confusion in the present example is a feeling one might have that antimilitary and pacifist are the same. Not eveybody agrees on this (not even the authors among themselves!) but if it helps, replace "antimilitary" by another predicate semantically unrelated to all the others.

C2-1. **Erik Sandewall**

In answer to Thielscher: but if unintended conclusions are obtained, then one ought to consider that the problem is with the axiomatisation and not with the logic itself.

C2-2. **Michael Thielscher**

(To Sandewall) Do you think it is the fault of the axiomatization in this case?

C2-3. **Erik Sandewall**

I would certainly *hope* so. More generally, since a variety of different strategies for multiple inheritance with exception has been studied in the specialized literature on that topic, one ought to require from a nonmonotonic logic that is used for this purpose that it be capable of expressing several of these inheritance strategies, and not just one of them.

C2-4. **Grigoris Antoniou** (2.2)

In a default logic setting sceptical reasoning is obtained by considering the intersection of extensions. One might say that defeasible logic takes a more intelligent approach by reasoning **along the way** what will and will not be part of this intersection, and makes use of this information accordingly.

Regarding the axiomatization of example 4.5, if one takes the view that antimilitary and pacifist are the same, then indeed they should be represented as one predicate. In that case we would get an undisputed result (nothing can be proven about pacifist/antimilitary).

Q3. **Erik Sandewall**

In response to your question during the workshop presentation about references to similar or related work in the literature, I suggest that there is an apparent relation to hypothesis logics, which have been developed by several authors. The difference is as follows: in your logic, tagged formulae using the operator *partial* appear in the conclusions, but not in the theories from which conclusions are obtained. In hypothesis logics, on the other hand, tagged formulae can be used everywhere, and therefore one does not need the additional types of implication that you use. Instead of formulae of the form *bird*(x) *typically-implies flies*(x), hypothesis logics use formulae of the form *bird*(x) implies *default*(*flies*(x)). (Exact references will be added).

Q4. **Alessandro Provetti**

(Question not recorded).

A4. **Grigoris Antoniou**

We believe that our type of theory, which uses rules plus a priority relation, is strictly more expressive than rules plus negation as failure. (References).

Q5. **Erik Sandewall**

I also have a question about your motivation for introducing a new nonmonotonic logic for the explicit purpose of commonsense reasoning. It seems to me that there is only a limited number of uses of nonmon reasoning, and that each of those application is being studied by separate methods and area-specific techniques. Multiple inheritance and reasoning about actions are two examples of such uses; there are of course a few more but not so many. Given this, wouldn't it be better to view nonmonotonic logics as instruments that are to be used for the area specific developments? Relating a nonmon logic directly to specific examples of commonsense reasoning seems to take in one step what would be better analyzed in two steps.

A5. **Grigoris Antoniou**

One intended use of our logic is for formulation of specifications of controllers, which are to be compiled in a chain of the following form: (controllers in defeasible logic) $\Rightarrow$ (tables) $\Rightarrow$ microchips.

C5-1. **Bob Kowalski**

(To Sandewall) There are plenty of other uses of nonmonotonic reasoning besides the ones you mentioned, in particular, for legal reasoning. (Agreement from McCarthy).

Q6. **Graham White**

This approach seems to get quite complicated. What does it take to write down a formal description of a domain? and why did you choose this kind of logic in the first place?

A6. **Grigoris Antoniou**

The original logic was not designed by us; it had been published before. Our goal is to try it out, and to relate it to other approaches in the area. We also see a need to develop systems for special cases to meet specific applications.

C6-1. **Pat Hayes**

In defense of the proposed approach, it seems to me that we are never going to get our axioms bug-free, and one way of dealing with that is to use systems like this one.

C6-2. **Graham White**

Still I wonder why it has to be so complicated.

C6-3. **Pat Hayes**

That seems to be unavoidable.

C6-4. **John McCarthy**

I don't think it's a question of bugs. (...)

C6-5. **Pat Hayes**

Maybe "bugs" was the wrong word; I'm thinking of generalizations that seem plausible when you write them but which are in fact false. Rather than replacing them, it may be better to leave them where they are, but to be prepared to add other facts that cancel them out.

Q7. **Bob Kowalski**

The presented article proposes a particular operational semantics. We now know that all of these kinds of systems (logic programming, etc) allow for multiple semantics... (End of the question not recorded).

A7. **Grigoris Antoniou** (2.2)

Indeed it is important to find other semantic accounts of the logic. We are particularly looking at argumentation-theoretic characterisations.

Vol. 2: *63-64*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/020/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Yves Moinard and Raymond Rolland

Circumscriptions from what they cannot do (preliminary report)

Presented at the workshop by Yves Moinard

Q1. **Pat Hayes**

What properties must have your preference relation, must it be transitive?

A1. **Yves Moinard** (27.1)

At this point (the question arrived at the begining of the exposition, corresponding to our prop. 4.2) we require no property at all: $<$ may be just any binary relation. People have already argued about the importance of requiring, or not requiring, transitivity from a knowledge representation perspective, and I have nothing significant to add to these discussions.

Whatever opinion you have on the subject, it is important to give the properties shared by any preferential entailment, without condition for the relation $<$. Indeed, if in some particular application, it happens that $<$ is transitive, and this fact is never used, then this shows that an important point has been missed somewhere: the study of the "reasoning properties" of the preferential entailment used here must be pursued.

Moreover, in our work, it happens later that we need transitivity. Indeed, we establish an exact correspondance between the relation $<$ over the set of models and the converse of the entailment relation between some subset of the set of the inaccessible formulas. As the entailment relation is transitive, we clearly need a transitive relation $<$. This corresponds roughly to the fact that we need cumulativity if we want to be able to express a preferential entailment from its inaccessible formulas.

Q2. **Tom Costello**

Are there any connections between your work and SCAN?

A2. **Yves Moinard** (27.1)

I do not think there are, but I do not know really. It seems to me unlikely that the formulas encountered while using SCAN have anything to do with the formulas in the set $X_<$ On the other hand, our work is clearly related to some aspects in Bossu and Siegel's method (AI journal, 1985) or Suchenek's method (Journal of Automated Reasoning, 1997).

Till now, we have only investigated the knowledge representation (or common sense reasoning, as we understand the expression) aspects. We would like to investigate in the future the possible interest of our work from a computational point of view.

Vol. 2: 65-77, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/042/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Peter Grünwald

Ramifications and sufficient causes

Q1. Erik Sandewall

You mentioned in your talk that Lin's use of the mnemonic "causes" tends to be misleading. I agree with this, and indeed the observation from our group to his work was that it was more or less a reformulation of what had already been done using occlusion. However, I wonder if there isn't a similar problem with your use of the mnemonic "do". When you characterize the *Toss* event using formulae such as $Do(Heads(1), true)$, then who is it that is doing something?

A1. Peter Grünwald (19.2)

Yes and no. 'Yes' in the sense that the word 'Do' may be misleading too - - indeed, nobody needs to be doing something when a coin falls on the table. 'No' in the sense that I am very explicit in my article about the 'physics' of Do: I state precisely what has to happen in a domain in order for $Do(X, B)$ to be the case in that domain; no such statement about *Caused* can be found in Lin's work.

Namely, $Do(X, B)$ means that an intervention takes place that sets the value of X to B. A more appropiate term might indeed be $Set(X, B)$ or $Intervention(X, B)$. I used Do to stay close to Pearl's notation, see [c-tark-96-51].

Lin's use of the predicate *caused* has not such a clear interpretation. It gets an implicit interpretation by the axioms it is involved in (like $Caused(X, s, True) \rightarrow Holds(X, s)$) but it is not a priori clear what has to be the case in the domain of interest in order for *Caused* to be true.

In other words, we do not really know what we are trying to model by the predicate *Caused* - that's the essential difference with *Do*.

References:

[c-tark-96-51] Judea Pearl. *Causation, Action, and Counterfactuals.* Proc. Theoretical Aspects of Rationality and Knowledge, 1996, pp. 51-73.

C1-1. **Erik Sandewall** (21.2)

Peter,
You wrote

> *... I am very explicit in my article about the 'physics' of Do: I state precisely what has to happen in a domain in order for $Do(X, B)$ to be the case in that domain; no such statement about Caused can be found in Lin's work.*
>
> *Namely, $Do(X, B)$ means that an intervention takes place that sets the value of X to B. ...*
>
> *Lin's use of the predicate caused has not such a clear interpretation.*

But Lin [c-ijcai-95-1985] wrote:

> *The ternary predicate Caused - for any fluent p, any truth value v, and any situation s, $Caused(p, v, s)$ is true if the fluent p is caused (by something unspecified) to have the truth value v in the situation s.*

What is the difference?
Erik

References:

[c-ijcai-95-1985] Fangzhen Lin. *Embracing Causality in Specifying the Indirect Effects of Actions.* Proc. International Joint Conference on Artificial Intelligence, 1995, pp. 1985-1991.

C1-2. **Peter Grünwald** (23.2)

The difference is that it is not clear what exactly has to happen (in the real world) for a fluent p 'to be caused to have truth value v'.

For example, $Holds(Sw1(Up), s) \wedge Holds(Sw2(Up), s) \rightarrow Caused(Open, s)$ is how Lin formalizes the suitcase domain. This translates to 'if both switches of the famous suitcase are up, then it is caused to be open'

But is the suitcase caused to be open by the mere fact that the switches are up? Shouldn't it rather be: 'the suitcase is caused to be open only if the the switches are *put* in the up position? One may argue about this - as one may argue about the appropiateness of lots and lots of uses of the word 'caused' in everyday language: in different situations the word can stand for different things. Because of this ambiguity, I argue, theories of action in which 'caused' is used as an atomic concept, *intended to be used whenever we use the word 'caused' in everyday language*, will run into trouble at some point.

I think the concept of 'intervention' suffers no such ambiguity. Suppose I give you a description of the complete state of a reasoning domain at time t (or in a situation s, if you like). Then you will usually be able to say whether or not an intervention takes place at that point in time/ in that situation. It will not always be clear whether or not something is 'caused to have truth value v' at that point in time/ in that situation. In other words, to put it very bluntly, 'I know what I'm modelling' whereas Lin does not.

C1-3. **Vladimir Lifschitz** (25.2)

In connection with Lin's suitcase example, Peter Grünwald asks:

> *Is the suitcase caused to be open by the mere fact that the switches are up? Shouldn't we rather say that the suitcase is caused to be open if the the switches are* put *in the up position?*

To support Lin's view, let me quote from an article by John Searle in the New York Review of Books (1995, Vol. 17, No. 42):

> *In our official theories of causation we typically suppose that all causal relations must be between discrete events ordered sequentially in time. For example, the shooting caused the death of the victim. Certainly, many cause and effect relations are like that, but by no means all. Look around you at the objects in your vicinity and think of the causal explanation of the fact that the table exerts pressure on the rug. This is explained by the force of gravity, but gravity is not an event. Or think of the solidity of the table. It is explained causally by the behavior of the molecules of which the table is composed. But the solidity of the table is not an extra event, it is just a feature of the table.*

- Vladimir

C1-4. **Fangzhen Lin** (25.2)

I was delighted and flattered, on the one hand, to learn that one of my previous work had been the focus of discussions in recent Newsletters on Reasoning about Actions and Changes. On the other hand, I was also disappointed and disturbed that much of the remarks about my work were either based on unsubstantiated claims (see below) or plain personal insults such as Peter Grünwald's indignity "'I know what I'm modelling' whereas Lin does not." Hey man, are you really that desperate that you could not find any other way to defend your work?

I don't think this is what this newsletter was set up for, and I'm sure this is definitely not what this community needs.

> *to be misleading. I agree with this, and indeed the observation from our group to his work was that it was more or less a reformulation of what had already been done using occlusion. However, I wonder if there isn't a*

Whoever made this claim, could you please substantiate it? By my understanding of "reformulation", I'm expecting at least something like:

1. A pointer to work on occlusion before IJCAI-95 that discusses examples similar to those in my IJCAI-95 paper [c-ijcai-95-1985].
2. A pointer to work on occlusion before IJCAI-95 that proves results similar to those in my IJCAI-95 paper.
3. A reformulation of my predicate $Caused(p, v, s)$ in the language of occlusion.

- Fangzhen Lin

References:

[c-ijcai-95-1985] Fangzhen Lin. *Embracing Causality in Specifying the Indirect Effects of Actions.* Proc. International Joint Conference on Artificial Intelligence, 1995, pp. 1985-1991.

C1-5. **Erik Sandewall** (25.2)

Peter, Vladimir, and Fangzhen

(The following combines answers to your respective comments to the discussion)

A common feature in much of the recent work is the use of an operator that "enables" change, that is, the operator excepts a feature from the requirement or the preference of persistence. The neutral term exception operator then covers both occlusion (often written as a predicate X), Vladimir's "release" operator, Lin's "Caused" operator, and Peter's "Do" operator, and also e.g. the negation of the "persistent" operator of del Val and Shoham. (An important early reference that articulates the use of such an operator in a general way is Hector Geffner's paper at KR 1989, [c-kr-89-137]).

Anyway, when such logics are used and one is to express that a certain "cause" has a certain "effect", then one has in principle three possibilities:

- The antecedent may refer to the current (new) value of some fluent or combination of fluents.
- The antecedent may refer to the fact that the value of some fluent has changed. For generality, we can view this as a "discontinuity" so that it allows both continuous-valued and discrete-valued fluents. Possibly the antecedent refers to a combination of current value(s) and current change(s) in fluents. (For continuous change, we also have the case where a continuously changing fluent has reached a certain value, that is, a discontinuity of equality).
- The antecedent may refer to the exception operator, possibly in combination with current fluent value(s). This is what happens if "Caused implies Caused".

All of these are instances of what Fangzhen calls "fluent-based causes", that is, causal relationships that only refer to values or changes of fluents, but not to the occurrence of actions.

Peter, you seem to argue that (3) is the right approach, and quote Fangzhen as using (1); Vladimir defends this with the reference to Searle. However, I observe that in his 1995 article, Fangzhen allows both (1) and (3). The general form of causal rules in his formula (16) allows one or more uses of the exception operator in the antecedent. However, example (22) does not use it, apparently in order to assure that the minimization will be computable by Clark completion. Isn't that a problem in your (Peter's) case as well?

Vladimir, I don't see why everything that's called "causation" in natural language has to be handled with the same logical devices. Seconding Peter's point of view, why can't we use one formalization for causation chains, where one change causes another change, and another formalization for static dependencies? The *causal propagation semantics* that I introduced in my KR 1996 paper [c-kr-96-99] and used for assessments of several approaches to ramification, is exactly a way of formalizing those cases where one change causes another. The kinds of situations that Searle refers to, or the early "dumbbell" example of Ginsberg and Smith represent arguably another class of phenomenona.

For causal propagation of change, there remains the choice between approaches (2) and (3) above: should the rule trigger on the fact that change actually occurred, or on the fact that change was enabled? This is important especially when actions are nondeterministic.

In the Features and fluents tradition, we first used alternative (3) (tech reports only), but then two group members, Tommy Persson and Lennart Staflin argued that (2) was the right way to go. Their paper at ECAI 1990, [c-ecai-90-497], uses (2) for characterizing indirect change. Their paper contains examples from continuous domains, but the formalism is also defined for, and applies equally well to the discrete case. (A revised version of the paper [r-linkoping-90-45] extends the work) In their approach, they wrote rules where the antecedent requires a specific fluent to actually have changed its value, but they also allowed to specify or to restrict the new or the old value.

We were not alone, of course. The paper by del Val and Shoham at IJCAI '93, [c-ijcai-93-732] defines a predicate $persistent(p, s)$ that is like the negation of occludes, and an entailment method that (I believe) does chronological minimization of unoccluded change. They mention the need for restrictions on the persistence axioms, but these are quite generous:

> *some further requirements are needed in order to ensure that the extension of persistent at any state depends only on the current and possibly past state of the database, not on future states...*

Also, I imagine that the event calculus is now able to represent these things, and it would be interesting to hear what the history was there (Murray? Rob?).

Finally, in order to answer Vladimir's question

> *Can you please explain this?*

because of my parenthetical remark to Peter, after having agreed with him that the use of the mnemonic "causes" may be misleading,

> ... the observation from our group to [Lin's] work was that it was more or less a reformulation of what had already been done using occlusion

and Fangzhen's question

> *By my understanding of "reformulation", I'm expecting at least something like:*
>
> 1. *A pointer to work on occlusion before IJCAI-95 that discusses examples similar to those in my IJCAI-95 paper.*
> 2. *A pointer to work on occlusion before IJCAI-95 that proves results similar to those in my IJCAI-95 paper.*
> 3. *A reformulation of my predicate Caused(p,v,s) in the language of occlusion.*

First, Fangzhen, if you felt I was being rough, I apologize, it was not my intention. I don't know if it's worthwhile to go so deeply into this, but since you both ask the question let me summarize the background. As several authors are now using operators similar to the family of occlude/ release/ persistent/ caused/ do/, a clarification of history may in fact be of some general interest. The following are the specific contributions of Fangzhen's 1995 paper [c-ijcai-95-1985] according to its abstract:

> *... we argue that normal state constraints that refer only to the truth values of fluents are not strong enough for [specifying the effects of actions using domain constraints], and that a notion of causation needs to be employed explicitly.*

See e.g. the Persson-Staflin paper [c-ecai-90-497]. Several papers in our group during 1988-1993 used a unique concept, "occlusion" or "explanation" of change for several purposes: for imprecise timing of changes within an action with extended duration, for nondeterminism, and also for "our intuition that a discontinuity should have a cause" (Persson/ Staflin). The abstract continues:

> *Technically, we introduce a new ternary predicate... Caused(p, v, s) if the proposition p is caused... to have the truth-value v in the situation s.*

Compare e.g. my article in Journal of Logic and Computation, 1994 [j-jlc-4-581], section 7.2 (misprint corrected):

> *Here $[s, t]p := F$ is an often-useful abbreviation for $(s, t]Xp \wedge [t]p = F$. Informally, it is read as saying that the feature p changes its value to become F some time during the interval $[s, t]$.*

This reduces to *Caused*(p, F, s) if $s = t$. Finally, the abstract says:

> *Using this predicate, we can represent* fluent-triggered *[causal statements]*

Compare the 1990 Persson-Staflin paper, section 5 for some examples.

Besides the material that is mentioned in the abstract, the paper also contains a section 4, which begins

> *The procedure we followed in solving the suitcase problem can be summarized as follows: ...*

and which then also defines a class of theories for which the Clark completion is enough to compute the intended conclusions. The KR'96 paper by Gustafsson and Doherty [c-kr-96-87] shows the striking similarity between this procedure and PMON. The procedure for reducing PMON to a first-order formulation was described by Doherty and Lukaszewicz in their 1994 ICTL paper [c-ictl-94-82], together with similar reductions of all other entailment methods that were assessed in "Features and Fluents" [mb-Sandewall-94]. The generalization of PMON to allowing fluent-based causal rules, is trivial. The reduction also generalizes for rules of types 1 and 2 above. I believe that Clark completion also has trouble with type 3?

We should of course not get enmeshed in priority debates, neither in this Newsletter nor anywhere else. I do feel however that many papers in our area get published with very incomplete accounts of previous and related work. This is unlikely to change because of the constraints that affect us all: limitations on our time, and on the number of pages allowed for each article. No reason to cast stones, therefore, but maybe the section on "related work" in research papers ought not to be our only mechanism for assembling topic-specific surveys and bibliographies, and possibly the present debate forum could serve as a complement. Additional contributions are invited to this account of recent history, therefore.

References:

[c-ecai-90-497] Tommy Persson and Lennart Staflin. *A Causation Theory for a Logic of Continuous Change.* Proc. European Conference on Artificial Intelligence, 1990, pp. 497-502.

[c-ictl-94-82] Patrick Doherty and Witold Lukaszewicz. *Circumscribing Features and Fluents.* Proc. International Conference on Temporal Logic, 1994, pp. 82-100.

[c-ijcai-93-732] Alvaro del Val and Yoav Shoham. *Deriving Properties of Belief Update from Theories of Action (II).* Proc. International Joint Conference on Artificial Intelligence, 1993, pp. 732-737.

[c-ijcai-95-1985] Fangzhen Lin. *Embracing Causality in Specifying the Indirect Effects of Actions.* Proc. International Joint Conference on Artificial Intelligence, 1995, pp. 1985-1991.

[c-kr-89-137] Hector Geffner. *Default Reasoning, Minimality and Coherence.* Proc. International Conf on Knowledge Representation and Reasoning, 1989, pp. 137-148.

[c-kr-96-87] Joakim Gustafsson and Patrick Doherty. *Embracing Occlusion in Specifying the Indirect Effects of Actions.* Proc. International Conf on Knowledge Representation and Reasoning, 1996, pp. 87-98.

[c-kr-96-99] Erik Sandewall. *Assessment of ramification methods that use static domain constraints.* Proc. International Conf on Knowledge Representation and Reasoning, 1996, pp. 99-110.

[j-jlc-4-581] Erik Sandewall. *The Range of Applicability of some Non-monotonic Logics for Strict Inertia.* Journal of Logic and Computation, vol. 4 (1994), pp. 581-615.

[mb-Sandewall-94] Erik Sandewall. *Features and Fluents. The Representation of Knowledge about Dynamical Systems.* Oxford University Press, 1994.

C1-6. **Peter Grünwald** (27.2)

Vladimir,

I agree with you that the concept of 'causality in terms of intervention only' is too narrow (and Searle expresses it quite well!); I just think that most of the examples in common sense reasoning we are currently dealing with are simple enough to be formalized in 'interventionistic' terms.

On the technical side, I do think there are some advantages in using the formalization one would obtain by translating the 'reverse switches domain' (example 5 in my paper) to the suitcase domain, i.e. using $(holds \wedge caused) \vee (caused \wedge holds) \rightarrow caused$ instead of $holds \wedge holds \rightarrow caused$.

These advantages are the same as those outlined for the switches domain in my common sense paper (see example 5).

C1-7. **Peter Grünwald** (27.2)

Fangzhen,

First of all, let me apologize. I am very sorry to have insulted you with my remarks. This was never my intention. Too eager to write a reply quickly, I wrote it without thinking and I didn't realize that it would be taken literally.

In fact, I find your approach one of the most intuitive I have come across; I think we both have quite similar intuitions about we are modelling which makes it all the more silly to write 'you don't know what you're modelling'. Here is what I really wanted to say; I hope I'll be able to formulate it in a better way now:

'I think it is clearer what kind of situation can and what kind of situation cannot be modeled with the help of the 'Do' predicate (which is intended to stand for interventions) than with the help of the 'Caused' predicate (which, if I understand you correctly, is intended to be used in those situations in which we would use the *word* Caused in natural language).'

Now that I'm at it, let me elaborate on this a little bit:

Though their roots are different ('Do' coming from Pearl's structural equation theory and 'Caused' coming from your work), it turns out that *formally* speaking, 'Do' is almost the same as your 'Caused' : they are defined in very similar ways which may even be equivalent (though this hasn't been proven). One of the aims of (the long version of) my paper is to point this out. That it may make sense to point it out can be seen from example 5 in the paper (the 'reversed switches domain', originally due to Sandewall and/or Doherty, I believe). The Do-predicate occurs there in the following axiom:

$$Do(Light(t), TRUE) \rightarrow Do(Switch(t), TRUE)$$

(from the fact that the light has been put on we conclude that the switch has been put in the on position)

If we describe to one another, in natural language, what is modelled by this axiom, then we would probably not use any causal terms (does putting on the light cause the switch to be on?) (see the paper for details). However, () can be modelled equally well by your 'caused'-predicate as with the 'do'-predicate used here. It follows that your 'caused'-predicate has broader applicability than it might seem, which I think is worth noting.

Peter

C1-8. **Peter Grünwald** (27.2)

Erik,

You wrote:

> *Peter, you seem to argue that (3) is the right approach, and quote Fangzhen as using (1); Vladimir defends this with the reference to Searle. However, I observe that in his 1995 article, Fangzhen allows*

both (1) and (3). The general form of causal rules in his formula (16) allows one or more uses of the exception operator in the antecedent. However, example (22) does not use it, apparently in order to assure that the minimization will be computable by Clark completion. Isn't that a problem in your (Peter's) case as well?

Indeed, Fangzhen allows both approaches (1) [*holds* $\rightarrow$ *caused*] and (3) [*caused* $\rightarrow$ *caused*]. My point is only (see my answer in C1-7 to Lin's question in C1-4)) that my approach can help you in deciding what domains should be modelled by (1) and what domains should be modelled by (3). I may indeed have trouble with efficiently computing the minimization - that wasn't my primary concern in the paper.

C1-9. **Vladimir Lifschitz** (27.2)

Erik,

I am puzzled by your comments about Lin's predicate *Caused* in ENRAC 25.2. In your view, it is an "exception operator," like occlusion. My understanding is very different. Exception operators are as old as formal default reasoning; McCarthy called his exception operator *Ab*, and its predecessor *End* can be found in your paper written back in the 1970s. Lin's predicate *Caused*, on the other hand, is a new and original idea.

You quote from Lin's abstract:

> *Technically, we introduce a new ternary predicate... Caused*(p, v, s) *if the proposition p is caused... to have the truth-value v in the situation s.*

Then you write:

> *Compare e.g. my article in Journal of Logic and Computation, vol. 4, no. 5, 1994, section 7.2 (misprint corrected):*
>
> > Here $[s, t]p := F$ is an often-useful abbreviation for $(s, t]Xp \wedge [t]p = F$. Informally, it is read as saying that the feature p changes its value to become F some time during the interval $[s, t]$.
>
> *This reduces to Caused*(p, F, s) *if* $s = t$*. Finally, the abstract says:*

Two points:

First, if $s = t$ then your informal reading turns into "p changes its value to become F some time during the interval $[t, t]$." How can p change its value during an interval consisting of one time instant?

Second, Lin's paper and yours use different temporal formalisms. To compare the two approaches, let's translate your formulas into the situation calculus. The counterpart of the formula

$$(s, t]Xp \wedge [t]p = F$$

in the situation calculus is

$$Ab(p, a, s) \wedge Value(p, Result(a, s)) = F$$

where Ab is the abnormality predicate from the commonsense law of inertia

$$\neg Ab(p, a, s) \rightarrow Value(p, Result(a, s)) = Value(p, s)$$

The fundamental difference between (1) and *Caused*(p, F, s) is that the former contains an action variable, and the latter doesn't.

Lin's IJCAI-95 paper was a major development that has not been fully appreciated so far by the nonmonotonic community.

–Vladimir

C1-10. **Erik Sandewall** (27.2)

Vladimir,

I thought I had found a previously unused term when I used "exception operator" as a generic for "occlude", "release", "Caused", etc. I don't think you would claim that all of those are trivially reducible to abnormality?

Re

> *First, if $s = t$ then your informal reading turns into "p changes its value to become F some time during the interval $[t, t]$." How can p change its value during an interval consisting of one time instant?*

a feature is said to change its value at time t iff its value at t differs from its value at $t - 1$. This is for the case of discrete time; for continuous time you have to use the left limit value at time t. Then, a feature is said to change its value in the interval $[s, t]$ iff it changes its value in some point in $(s, t]$ (omitting the left endpoint of the interval), that is, either between s and $s + 1$, or..., or between $t - 1$ and t.

When I referred to the special case of $s = t$ in $[s, t]p := F$, I chose naturalness over formal precision, since it should have said $s = t - 1$. I thought introducing this technicality would have been an unnecessary detour.

Re

> *The fundamental difference between (1) and Caused(p, F, s) is that the former contains an action variable, and the latter doesn't.*

evidently this action variable is an artifact of your translation, since it is not present in the original formulation. Also, since Lin's paper does not use any *Ab* predicate (he minimizes *Caused* directly), translating my formulation to one using *Ab* doesn't particularly help the comparison. In fact, if you translate Lin's notation in the same way as you did for mine, you will have to introduce an action variable there as well since your *Ab* predicate is *defined* with that as its second argument. However, this has nothing to do with either Lin's formalism or mine.

The translation task is actually very simple, since my formalism allows for both linear time and branching time (F&F page 138 ff). For branching time, the notation $[s, t]$ is defined only if t is a direct or indirect successor of s, and it then refers to the path from s to t. Defining t^- as $t - 1$ in the case of integer time, and through the axiom $do(a, s)^- = s$ for the case of gof-sitcalc situations, the obvious translation of

$$(t^-, t]Xp \wedge [t]p = F$$

into Lin's notation is

$$\mathit{Caused}(p, F, t)$$

(Remember that $[t]p = v$ is the same as $H(p, v, t)$, and that this is a modified kind of equality, whose character can't be rendered in HTML). Both the informal explanation and the formal treatment of (1) and (2) are identical in the two approaches, for mainstream examples. Formula (2) is more compact, which is why I introduced the abbreviation

$$[t^{-}, t]p := F$$

Then it's clear that the difference is purely notational. If you don't want to represent the duration of actions, then you can of course modify the abbreviation so that only one timepoint is mentioned.

C1-11. **Fangzhen Lin** (27.2)

Erik,

Some thoughts about your note C1-5.

1. $Caused \not\equiv Occlude$. For example, $Caused(p, true, s) \rightarrow H(p, s)$. But there are no such axioms for $Occlude$. Notice that these axioms go a long way: in deriving a generic successor state axiom, new causal rules from old ones, making causal rules strictly stronger than corresponding state constraints...

2.

$$Caused(p, true, s) \not\equiv Occlude(p, s, s) \wedge H(p, s)$$
$$Caused(p, true, s) \not\equiv Occlude(p, s) \wedge H(p, s)$$

See (Gustafsson and Doherty, KR'96).

3. Regarding the previous results of occlude that you mentioned:

> ... we argue that normal state constraints that refer only to the truth values of fluents are not strong enough for [specifying the effects of actions using domain constraints], and that a notion of causation needs to be employed explicitly.

> *See e.g. the Persson-Staflin paper at ECAI 1990. Several papers in our group during 1988-1993 used a unique concept, "occlusion" or "explanation" of change for several purposes: for imprecise timing of changes within an action with extended duration, for nondeterminism, and also for "our intuition that a discontinuity should have a cause" (Persson/ Staflin). The abstract continues:*

Observe that none of them were about domain (state) constraints and the ramification problem. As you can see from the paragraph that you cited from (Lin 95), the main contribution of my paper has to do with constraints and in particular causal constraints. If I'm not mistaken, the main impacts of the set of quite related papers on actions in IJCAI'95 have to do with the ramification problem. They shed some new lights on the nature of this problem that had never before been thought of, and have just begun to bear fruits: for example, the work by Michael Thielsher on diagnoses that was discussed in this newsletter before, and the work on planning from causal theories using GSAT by Norm McCain and Hudson Turner (in KR'98).

4. Regarding the connection between your PMON and the minimization strategy that I used, yes, I'd be the first to admit that they are very similar. I apologize for not knowing it earlier. Thinking back, it's not surprising how they come together. As you know, I have been a devotee of what you call "filtered preferential entailment" approach for as long as I can remember,

going back when Yoav and I were working on our provably correct theories of actions.

The idea of minimizing *Caused* with *Holds* fixed was also straightforward as it was the obvious one that would yield the Clark completion of causal rules.

5. I enjoyed reading Gustafsson and Doherty's KR'96 paper. I learned a lot from that paper. Patrick and I had some correspondences on the related topics. I believe we found some common grounds, but none in terms of "one can be reformulated in terms of another". (Patrick, correct me if I'm wrong.) If anything, I wouldn't call PMON(RCs) a "trivial extension of PMON to causal rules", especially when *Occlude* is applied to a time point in PMON(RCs) while the same predicate is applied to a time interval in PMON. Many things are trivial in retrospect. (The idea of GSAT is so simple that Russell and Norvig even include it as an assignment in their AI textbook.) But in any event, this paper is more or less irrelevant as far as the history is concerned.

Last but not the least, Erik, I have all the respects and regards in the world for you and your work! And Peter, there is no hard feeling whatsoever.

Q2. **Camilla Schwind** (21.2)

As far as I have understood your approach, you do not use a causality predicate (or connective), but rather you define a causality relationship or causal implication by means of material implication as $Do(p, b) \rightarrow Do(q, b')$ or $a \rightarrow Do(p, b)$ as in example 3

(6) $Do(Alive(t), false) \rightarrow Do(Walking(t), false)$

and

(3) $Shoot \rightarrow Do(Alive(1), FALSE)$

But then you can derive classically, $A \wedge Do(p, b) \rightarrow Do(q, b')$, for any formula A. In other words, your causality relation appears to be **monotonic**. I think however that causal implication should be **non-monotonic**. For example, from you axiom (3), you can derive

$Shoot \wedge \neg - loaded \rightarrow Do(Alive(1), False)$

Does you approach behave in this way and do you think that causality should rather be non-monotonic?

A2. **Peter Grünwald** (23.2)

Yes, my 'causal implication' is monotonic. To handle examples like

$$Shoot \wedge \neg - loaded \rightarrow Do(Alive(1), False)$$

I would remodel my domain as follows:

$$Shoot \wedge -Abnormal_i \rightarrow Do(Alive(1), False)$$

All the nonmonotonicity is then put in the $Abnormal_i$ propositional variable. This is discussed, by the way, in the technical report INS-R9709 (basically a long version of my Common Sense paper) available at my web-page. For brevity, I excluded all talk of abnormalities from the Common Sense paper.

C3-1. **Vladimir Lifschitz** (21.2)

Hi Erik,

> *You mentioned in your talk that Lin's use of the mnemonic "causes" tends to be misleading. I agree with this, and indeed the observation from our group to his work was that it was more or less a reformulation of what had already been done using occlusion.*

Can you please explain this?
Regards, Vladimir

C3-2. **Erik Sandewall** (25.2)

See comment C1-5 above.

Electronic News Journal on Reasoning about Actions and Change
Vol. 2: *78-79*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/056/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Eyal Amir

Point-Sensitive Circumscription

Presented at the workshop by Tom Costello

Q1. Erik Sandewall

I have questions about your section 6, "Applying Point-Sensitive Circumscription to Theories of Action", which I can not make sense of. You first review the formulation of PCM (Prototypical Chronological Minimization) in pointwise circumscription, and then write (p. 71):

> *We do not claim that Point-Sensitive Circumscription can formalize the entire $\mathcal{K}$-**IA** ontological class, but that it can take you farther than Pointwise circumscription can.*

Let me briefly review the relevant aspect of Features and fluents, in order to put what follows into context. The $\mathcal{K}$-**IA** ontological class allows for scenarios with nondeterministic actions, actions with extended duration in time and internal structure, and all combinations of prediction and postdiction. It does not model causality or other forms of ramification, nor concurrency, etc.

For this $\mathcal{K}$-**IA** class, we have studied a number of *entailment methods*, each of which can then be reexpressed in e.g. pointwise circumscription. PCM is one of those entailment methods. Each of them is defined in terms of preference criteria or other selection criteria on models, and it is these criteria that may be translated e.g. to circumscription as a way of proceeding towards an implementation. For each entailment method, there is also an assessment of range of applicability. This assessment is model-theoretic, and is independent of the choice of implementation method.

In this framework it is perfectly reasonable, as you do, to evaluate a variant of circumscription with respect to whether or not it is capable of expressing a particular entailment method, or even a collection of them. This makes much more sense than analyzing circumscription variants directly in terms of test examples, since then so much depends on the choice of axiomatization.

Though I agree with your methodology up to that point, I do not understand the concrete case where it is applied. You only discuss how PCM is rendered in variants of circumscription, but Doherty has already shown how to express PCM in pointwise circumscription. If your variant of circumscription is able to handle more examples correctly within the range of $\mathcal{K}$-**IA**, it could therefore only be because it were an inaccurate rendering of PCM.

Also, although the range of applicability of PCM is only a limited part of $\mathcal{K}$-**IA**, there are other entailment methods such as PMON which are correct for $\mathcal{K}$-**IA**, and which can also be expressed in pointwise circumscription following Doherty. This makes it even more strange how point-sensitive circumscription can do something that pointwise circumscription can not, with respect to the $\mathcal{K}$-**IA** class.

You also discuss a simple case of ramification using the duality between the *dead* and *alive* properties. This of course is not within $\mathcal{K}$-**IA**, as you also observe. However, there are certainly known methods that handle such simple ramifications (as well somewhat less trivial ones) correctly and that have been expressed or re-expressed using pointwise circumscription.

Could you help me understand your contribution in these respects?

Vol. 2: *80-82*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/075/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Fausto Giunchiglia and Chiara Ghidini

Local Model Semantics, or Contextual Reasoning = Locality + Compatibility

Presented at the workshop by Chiara Ghidini

Q1. Tom Costello

You define your semantics in such a way that each viewpoint has a separate interpretation. I would have imagined using a single interpretation to which all the viewpoints can be related. In your six-box example, there would have been one box with six compartments for explicating both views.

A1. Chiara Ghidini

Usually we don't have the complete description of the world.

C1-1. Tom Costello

But a semantics is not supposed to work like that. It need not identify a single model; it defines a structure, and a given set of axioms may have a very large set of models, reflecting our lack of knowledge.

C1-2. Pat Hayes

I agree with Tom (Costello); the point is that there is only one reality.

Q2. Erik Sandewall

Regarding question 1, I want to support Chiara's approach. If we were to follow Tom's (Costello) advise, we would have to define how the unique interpretation is projected into the different views. That is not necessarily any simpler than using the compatibility relations.

My question regards what is the syntactic counterpart of the compatibility relations. It would seem that a given set of bridge rules corresponds to certain compatibility relations, and vice versa. Is this correct, and if so, can you characterize the relationship between bridge rules and compatibility relations?

A2. **Chiara Ghidini**

Yes, for limited cases of the bridge rules we can do that.

Q3. **Dov M. Gabbay**

I also want to defend Chiara's approach. This corresponds to the fibering of systems (...)

C3-1. **Pat Hayes**

(Interrupting Gabbay after a while): What in the world are you talking about? (...)

C3-2. **Dov M. Gabbay**

The thing is that you and the speaker don't understand each other...

C3-3. **Pat Hayes**

I don't understand you either. (...)

C3-4. **Dov M. Gabbay**

The whole thing is about mixing two languages. You evaluate one until you reach an expression in the other, then you transfer to the interpretation of the other language (...)

C3-5. **Graham White**

It would be interesting to know *when* it is possible to derive a global model from the compatibility relation.

Q4. **John McCarthy**

In support of Chiara's approach, I also don't believe it's possible to assume a single representation of the world. Given my representation of Pat's beliefs and Dov's representation of Pat's beliefs, is there underlying that a *true* account of Pat's beliefs?

Q5. **Tom Costello**

Comparing the two approaches of [c-fcs-98-114] (Pirri) and [c-fcs-98-75] (Ghidini), it seems that they are addressing similar topics but using opposite methods. Fiora (Pirri) decomposes a large thing into its parts, Chiara (Ghidini) takes local information and puts it together. The same words are being used for things that work in opposite directions.

References:

[c-fcs-98-114] Gianni Amati and Fiora Pirri. *Contexts as relativized definitions: a formalization via fixed points.* Proc. Formalization of Commonsense Reasoning, 1998, pp. 114-125.

[c-fcs-98-75] Fausto Giunchiglia and Chiara Ghidini. *Local Model Semantics, or Contextual Reasoning = Locality + Compatibility.* Proc. Formalization of Commonsense Reasoning, 1998, pp. 75-86.

C5-1. **Fiora Pirri**

The only difference is that Chiara can use two different languages; I just use one language.

Q6. **John McCarthy**

Chiara started out by saying "a context is...", but her definition seemed much more narrow than what was in my original paper about contexts. This is a general complaint I have about this kind of work - it is rarely concerned with lifting a formula from one context to another. There's a lot of other things that need doing as well.

C6-1. **Fiora Pirri**

I agree, but there is a problem with interpreting terms in contexts. A question to John again: how narrow is a context?

C6-2. **John McCarthy**

If I make a definition of the form "A context is..." then I would give a definition that's more narrow than what I'd be prepared to live with.

C6-3. **Pat Hayes**

At a recent workshop on context I collected definitions of what it is. They ranged from "a cultural viewpoint", via "a first-order theory" to "a point in spacetime". It seems to be difficult to get any agreement on that.

Q7. **Vladimir Lifschitz**

It is interesting to compare this paper with John's two approximations (?). Here there are no approximations. An interesting alternative would be to consider two observers, both of which see everything but with observation errors. Do you envisage considering this case?

A7. **Chiara Ghidini**

Yes, in federated databases there are examples of that.

Electronic News Journal on Reasoning about Actions and Change 83
Vol. 2: *83-83*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/087/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

François Lévy and J. Joachim Quantz

Representing Beliefs in a Situated Event Calculus

Presented at the workshop by J. Joachim Quantz

Q1. General discussion (7.1)

What you call "situations" in your approach does not appear to be similar to situations as used in situation calculus. So what are they? Can they be understood as "histories", so that $\mathit{Holds\text{-}at}\,(p, t, s)$ should be read as saying that p is true at time t in the history s?

A1. J. Joachim Quantz

Yes, but notice that the set of situations or histories is branching, so that two situations/ histories may be identical up to a point in time, and differ from there on.

Q2. Erik Sandewall

I like this; it allows us to reason about alternative futures with a moderate extension of the previous machinery. However, one obvious limitation is that beliefs can only be expressed relative to objective timepoints. There is no way of representing believed time. Do you have any suggestion about how the proposed framework can be extended so as to make this possible?

A2. J. Joachim Quantz

This is a topic for our further research.

Vol. 2: *84-85*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/099/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Graham White

Balls and Strings: Simulations and Theories

Q1. **Michael Thielscher**

We have done similar things in first-order logic. What is the advantage of switching to linear logic?

A1. **Graham White**

It provides a more elegant meta-theory.

C1-1. **Erik Sandewall**

What is the advantage of having a more elegant meta-theory?

C1-2. **Graham White**

It provides more elaboration tolerance and more modularity.

C1-3. **A workshop participant**

Is any publication available which demonstrates those advantages?

C1-4. **Graham White**

No, not yet.

Q2. **François Lévy**

If I were convinced by your arguments and wanted to use linear logic for an application, how do I use the connectives of linear logic to describe the world?

A2. **Graham White**

It is very simple, $\otimes$ for forming situations and $\multimap$ for rewrites.

Q3. **Pat Hayes**

What about if you wish to reason in different directions, and not only in the direction of the simulation, does linear logic do that?

A3. **Graham White**

Yes. (Details not recorded).

Q4. **Michael Thielscher**

This is important, I want to make sure I get it right. Is it the case that I can do reasoning with incomplete information about the initial state and reasoning backwards in time, using the same axioms describing the application and the same deductive procedure, if I use linear logic?

A4. **Graham White**

Yes, using Girard's fixpoint theorem you can do it.

C4-1. **Michael Thielscher**

Is that standard linear logic?

C4-2. **Graham White**

Well, it's second order linear logic.

C4-3. **Michael Thielscher**

What is the advantage of second order linear logic over first order predicate logic?

C4-4. **Graham White**

I prefer the semantics.

C4-5. **Pat Hayes**

What *is* the semantics?

C4-6. **Graham White**

It is based on sets of states in Petri nets. (The discussion continued).

Electronic News Journal on Reasoning about Actions and Change 86
Vol. 2: *86-86*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/114/

Protocol of the discussion at the FCS-98 workshop and of
the subsequent on-line discussion during February, 1998
about the following research article:

Gianni Amati and Fiora Pirri

Contexts as relativized definitions: a formalization via fixed points

Presented at the workshop by Fiora Pirri

Q2. John McCarthy

I would have thought that representing contexts as fixpoints would give you more problems than by other methods, for example as explicit objects. Can you compute the fixpoint?

A2. Fiora Pirri

Yes, of course.

Q3. Tom Costello

Your argument that contexts are implicit is based on them not being mentioned in natural language. However, that argument is irrelevant, and if accepted it could be applied against very many constructs in current work in our field. Wouldn't it be better to represent contexts by amalgamation, starting with small ones and composing larger ones from them?

A3. Fiora Pirri

One would then have the problem of how to design those elementary contexts in order that they do the right thing when composed. We do not have a methodology for finding out the right level of refinement.

Q5. François Lévy

What happens when viewpoints are inconsistent?

A5. Fiora Pirri

That is OK, they do not have to be consistent.

Vol. 2: *87-87*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/126/

Protocol of the discussion at the FCS-98 workshop and of
the subsequent on-line discussion during February, 1998
about the following research article:

Emilio Remolina and Benjamin Kuipers

Towards a formalization of the Spatial Semantic Hierarchy

Presented at the workshop by Emilio Remolina

Q1. Murray Shanahan

A general account of space should deal with both large-scale and small-scale shapes in an integrated fashion. Is your system really able to do that? Your use of the term "T-shape" suggests that the difficulty of maintaining a strict difference between these.

A1. Emilio Remolina

This is an open problem.

Q2. Erik Sandewall

How would it change your problem or its solution if the robot were to have a position sensor that is able to indicate its position with the full accuracy that is needed for distinguishing between several places that look similar, such as several corridor crossings? Current GPS technology would be able to provide this.

A2. Emilio Remolina

This would facilitate a lot, but our problem would basically remain the same.

Electronic News Journal on Reasoning about Actions and Change 88
Vol. 2: *88-88*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/138/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Philippe Balbiani and Luis Fariñas del Cerro

A relational model of movement

Presented at the workshop by Luis Fariñas del Cerro

Q2. Erik Sandewall

You are using a geometrically oriented approach. Another approach which has been used by a number of authors, including Shanahan and myself, is to embed differential equations within a first-order logic. This requires a multi-sorted logic with continuous time as one of the sorts; from there on it is a fairly straightforward extension of narrative time-line approaches that use discrete time, which means that it has the advantage or at least the potential of integrating well with solutions to other issues, such as inertia (generalizes to minimization of discontinuities), causality, and concurrency.

Question, therefore: What advantage do you see in the geometrical approach?

A2. Luis Fariñas del Cerro

I don't think logic has anything to do with continous change.

C2-1. Erik Sandewall

Certainly it is technically possible to have a logic where the real numbers is one of the sorts, and that's also what we have done. Given that both approaches exist, we ought to be able to compare them.

C2-2. Luis Fariñas del Cerro

(Additional discussion followed)

Vol. 2: *89-90*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/146/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Sven Ove Hansson and Renata Wasserman

Local Change: A preliminary report

Presented at the workshop by Renata Wasserman

Q1. John McCarthy

It seems to me that if you used formalized contexts, then you would be able to isolate different sets of propositions and say something about them in logic. In that way you wouldn't have to invent a new logic for this purpose.

A1. Renata Wasserman

The difference is not all that great.

C1-1. Pat Hayes

Yes, but if you reify formulae and quantify over them, then it *is* a new logic. (Agreement from McCarthy).

Q2. John McCarthy

A compartment could be a context, couldn't it?

A2. Renata Wasserman

We preferred to use the word "compartment" over "context".

Q3. Mikhail Soutchanski

(Question not recorded)

A3. Renata Wasserman

Another thing that we'd like to do is to facilitate retrieving the compartment. Predefined dependency relations between previous beliefs may be instrumental for that purpose.

Q4. **Tom Costello**

(Beginning of question not recorded). ... I understand your logic treats conjunction differently from having the two conjuncts separately. Isn't this very strange?

A4. **Renata Wasserman**

That depends on what you use for retrieving the relevance.

Q5. **Bob Kowalski**

I was going to propose using metalogic, but then the problem is that you are naming sentences rather than propositions. A possible way out may be to use metalogic, but in such a way that it operates on a canonical form for the sentences, for example, clausal form.

A5. **Renata Wasserman**

I don't think that we (people) represent sentences in canonical form.

C5-1. **Pat Hayes**

But that's not incompatible.

C5-2. **Graham White**

(Referring to the discussions on both questions 4 and 5) But there are situations where different representations *are* treated differently.

C5-3. **Pat Hayes**

(To Graham White) Possibly, but that opens a can of worms. Imitating nonlogical behavior will carry too far.

Anyway, these logics *have* been studied argumentation theoretically, so there is plenty to choose from. If you don't like one semantics, then go back and pick another one.

C5-4. **Erik Sandewall**

(To Hayes) That's dangerous advise. Changing semantics is like throwing a big switch on your whole system, with all its ways of representing different aspects of the application at hand. Unless, of course, you are able to interpret different *parts* of the system with different semantics - but that's a big can of worms as well.

C5-5. **Tom Costello**

This all reminds me of the fights between programming languages, and the experience is that the worst language always wins.

Vol. 2: *91-91*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/166/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Frans Voorbraak

A Nonmonotonic Observation Logic

Q1. A workshop participant

You propose using default logic as the basis of your approach, but then you do not seem to make any concrete use of it. Why do you need it at all?

A1. Frans Voorbraak

It is not used in a concrete way now, but it will in the forthcoming extensions where we allow for defeasible background knowledge as well as defeasible observations.

Vol. 2: 92-93, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/183/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Brian Knight, Taoxin Peng, and Jixin Ma

Reasoning about Change over Time: Actions, Events, and their Effects

Q1. Erik Sandewall

A fairly large part of your work is devoted to the time structure. Then, under "future work" you mention addressing the topic of continuous change. Won't you then have to use the real numbers (or something similar) as time structure?

A1. Jixin Ma (26.1)

Yes, though it is not necessary to directly use the real numbers as the time structure. In fact, a general time theory based on both points and intervals is proposed previously [j-cj-37-114]. With the corresponding density axioms, it seems appropriate for addressing continuous change.

References:

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

C1-1. Erik Sandewall

In this case, why don't you introduce them at once, even for the present work? Is there some essential aspect of the present results that you are going to lose when you proceed to continuous change?

C1-2. Jixin Ma (26.1)

First of all, the approch addressing both moments and points as primitive is proposed in the paper in order to overcome the so-called Dividing Instant Problem, i.e., the problem in specifying if intervals is "closed"/"open" at their ending points in the case intervals are constructed out of points. It is argued again and again by many researchers that no solution is satisfactory to the question of what preconception about the closed/open nature of intervals should be taken:

- If all intervals include their ending-points, then adjacent intervals would have ending-points in common. Hence, if two adjacent intervals correspond to a given fluent and its negation, respectively, there will be a point at which the fluent is both true and false.
- Similarly, if all intervals don't include their ending-points, there will be points at which the truth or falsity of some fluents are undefined.
- Another approach is to take point-based intervals as semi-open (e.g., all intervals include their left ending-points, and exclude their right ones) so that they may sit conveniently next to one another. However, on the one hand, since this approach insists that every interval contains only a single ending-point, the choice of which ending-point of intervals should be included/excluded seems arbitrary, and hence unjustifiable and artificial. On the other hand, although the approach may offer a solution to some practical questions, there are some other critical questions which remain problematical. The fundamental reason is that in a system where time intervals are all taken as semi-open, it will be difficult to represent time points in a consistent structure so that they can stand between intervals conveniently.

There are three reasons of taking a discrete time structure:

1. It is simple and adequate for many practical applications.

2. It is noted that, if time intervals are addressed as temporal objects which can be associated with fluents by meta-predicates such as "HOLDS" or "TRUE", one shall face the possibility that a fluent might be neither true nor false throughout some interval. In the worst case, the so-called intermingling problem may arise, that is the possibility of a fluent changing its truth value infinitely often in an interval with a finite duration. This will lead to some difficulties in characterising the relationships between the negation of a fluent and the negation of involved sentances.

The approach that addresses intervals as a finite ordered union of points/moments sucessfully overcomes (or at least bypasses) such a problem. For the purpose of modelling continuous change, it seems necessary (this maybe arguable) to extend such a discrete time model to a dense one. However, in this case, some careful treatments have to be taken for dealing with the Intermingling Problem.

3. Since times are explicitly addressed in the formalism, the corresponding Frame Problem would become more difficult to be handled if the time structure is modelled as dense in the first place. Hence, the strategy is to tackle the simple case (that is, the discrete model) first, then to extend the idea to the general case.

Electronic News Journal on Reasoning about Actions and Change 94
Vol. 2: *94-94*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/217/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Enrico Giunchiglia and Vladimir Lifschitz

An Action Language Based on Causal Explanation: Preliminary Report

Presented at the workshop by Vladimir Lifschitz

Q1. Erik Sandewall

The paper is a preliminary report of a new action language. What do you expect to achieve by introducing this language which has not already been achieved by existing work on expressing causation in other approaches, such as the event calculus or the TAL languages of Doherty et al?

Q1. Vladimir Lifschitz (5.2)

The action language $\mathcal{C}$ (as well as the McCain/Turner causal logic that it is closely related to) may be of interest because of the simplicity of its syntax and semantics. Although it has only two kinds of propositions – static and dynamic laws – the language is quite expressive, as the examples in the paper demonstrate. The semantics of these propositions is defined in terms of simple and well-understood objects-transition diagrams. And the definition of the semantics in terms of classical logic is just a half-page long.

Second, the view of inertia in $\mathcal{C}$ is different from what we see in most other work on representing actions. Actions take place against the backdrop of the course of nature. In the past, the focus on inertia has emphasized worlds in which the course of nature is static. Our pendulum example illustrates the fact that inertia-like laws can describe the dynamic course of nature, and they can be overriden by effects described in axioms about actions, just like inertia. In the language $\mathcal{C}$, both inertia and the inertia-like pendulum law are treated syntactically in similar ways – as dynamic laws that do not include action names.

Vol. 2: *95-96*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/239/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Michael Thielscher

Steady Versus Stabilizing State Constraints

Q1. Erik Sandewall

You make a distinction between steady and stabilizing constraints, and characterize both using causal relationships. However, is it really appropriate to use the term "causation" for the results of the steady constraints?

(Answer not recorded).

Q2. Erik Sandewall

With respect to the stabilizing constraints, the important thing about them seems to be that they involve a time delay. If this is so, would it not be more natural to use an explicit timeline and to position the successive changes along it, using short time delays?

A2. Michael Thielscher

The question whether causes must truly precede effects rather than effects just not having to precede causes, is among the many disputed general problems of causality. The classical example in philosophical literature is this: Can we say that the fact that Socrates died caused Xanthippe to become a widow? My answer to that question is positive, and since I'm not alone I feel legitimated to call "causal" the relationship between these two events. (Clearly, the underlying state constraint in this example is a steady one.)

With respect to the introduction of explicit time to model stabilizing indirect effects: In so doing one lowers the level of abstraction–which is unnecessary and hence contradicts the commonsense principle of choosing the highest level of abstraction which is just sufficiently low to talk and reason about the domain in question. Everyone who ever in one of his or her papers has considered an example involving a stabilizing state constraint but has not used delayed effects to model this example, is a (maybe unintentional) witness for the case.

Q3. **Peter Grünwald**

How would you handle the suitcase example in your approach?

(Answer not recorded).

Vol. 2: 97-97, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/249/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Tom Costello

Minimizing the Effects of Actions

Q1. Erik Sandewall

You were saying that explanation closure requires determinism, and this seems to be an important point in the motivation of your approach. I do not agree with this point; with the right formulations it is possible to use explanation closure even for the action laws of nondeterministic actions.

Electronic News Journal on Reasoning about Actions and Change 98
Vol. 2: *98-98*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/263/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Tom Costello and Anna Patterson

Quantifiers and Operations on Modalities and Contexts

Presented at the workshop by Tom Costello

Q1. A workshop participant

This seems to be quite similar to dynamic logic. How does it differ from it and advance over it?

A1. Tom Costello

The basic operations are the same, but here we can quantify over modalities.

Vol. 2: *99-101*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/277/

Protocol of the discussion at the FCS-98 workshop and of
the subsequent on-line discussion during February, 1998
about the following research article:

Giuseppe De Giacomo, Ray Reiter, and Mikhail Soutchanski

Execution Monitoring of High-Level Robot Programs

Presented at the workshop by Mikhail Soutchanski

Q2. Rob Miller

You assume that the robot is able to sense all "actions". That is not realistic, is it?

A2. Mikhail Soutchanski

Originally, we defined the logic so that it sensed values of fluents, but this was technically more difficult to deal with. The present account is technically better.

C2-1. Rob Miller

But assuming that the values of fluents can always be correctly sensed is also not realistic. Observations will sometimes be faulty.

C2-2. Mikhail Soutchanski

Agreed.

Q3. A workshop participant

Doesn't your account indicate a basic design flaw in Golog, namely that it ought to have had an exception mechanism?

A3. Mikhail Soutchanski

(Answer not recorded).

Q4. **Javier Pinto**

How do you consider using the notion of relevance for prediction, in order to cut down the space of what has to be predicted in a given failure situation?

A4. **Mikhail Soutchanski**

We have thought about it, but we have no answer yet.

Q5. **Alessandro Provetti**

Since there is no notion of backtracking, why did you decide to use Prolog for the implementation?

A5. **Mikhail Soutchanski**

Its properties are used for the cautious interpreter.

Q6. **Murray Shanahan**

When an exogenous "action" has occurred, the system invokes the planner in order to find a way of getting to a state where the plan can be continued. This seems to require it to know about the preconditions for each action or action sequence, and not merely the postconditions of the whole program.

A6. **Mikhail Soutchanski**

No, the postconditions are enough.

Q7. **A workshop participant**

What if someone else comes in and completes building the tower while the robot is building one, will it then knock it down in order to be able to build it itself?

A7. **Mikhail Soutchanski**

It is a research issue how to avoid having to do that.

C7-1. **Erik Sandewall**

How can it possibly be anything but trivial? If you are going to pursue the existing program, you must find a plan for getting to a state where it can be continued. If you are going to shortcut it and find an alternative plan directly to the desired goal of the given program, then you have to find a plan for getting to a state that's characterized by the program's postconditions. If you have a system that does one of these things, it should be able to do the other. It should also be able to try both alternatives, compare the plans, and pick the best one.

C7-2. **Mikhail Soutchanski**

Programs may be fairly long, and we do not wish to engage the planner for making long plans, only short ones, which is usually sufficient for recovering to a state where the program can be continued.

C7-3. **Michael Thielscher**

(Referring to question 6) But if one does not look for going directly to the program's goal, but only for recovering, then it seems that preconditions of remaining actions *are* necessary, and the postcondition of the entire program is irrelevant.

Q8. **François Lévy**

You do not seem to rely on work in robotics or in planning. Why is that so?

A8. **Mikhail Soutchanski**

The work in robotics and in algorithmic planning assumes linear plans, but we need to have loops. Work in reactive planning doesn't take context into account.

Electronic News Journal on Reasoning about Actions and Change 102
Vol. 2: *102-102*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/298/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Simon Parsons, Carles Sierra, and Nick R. Jennings

Multi-context argumentative agents

Q1. Tom Costello

Do you have any formal results on whether or when your system is consistent?

A1. The Speaker

The beliefs of different agents do not have to be consistent.

C1-1. Tom Costello

That was not my point; the point was whether the overall system which characterizes and reasons about all the agents is consistent. Consistency is an absolutely essential property of such systems; without it they are useless.

(Answer not recorded).

Q2. Erik Sandewall

Does the account of the inference performed by each of the agents allow any nonmonotonicity, as required e.g. for prediction, or for default beliefs?

A2. The Speaker

Prediction: not yet. Default beliefs are obtained by context transfer.

Q3. Pat Hayes

Can the agent believe $p \wedge \mathsf{B}\, \neg p$?

A3. The Speaker

In the present system, it does not notice that problem.

Vol. 2: *103-104*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/319/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

John Bell and Zhisheng Huang

Seeing is Believing

Presented at the workshop by John Bell

Q1. Several participants

Your formula (3) must be completely wrong. It says that if α is not perceived, and if it is neither believed nor disbelieved, then it is both believed and disbelieved.

A1. John Bell

In this case we have:

$$P(a)(t) \wedge \neg \mathsf{B}\ (\neg a)(t) \rightarrow P(a)(t+1)$$
$$\neg P(a)(t) \wedge \neg \mathsf{B}\ (a)(t) \rightarrow \mathsf{B}\ (\neg a)(t+1)$$

Suppose that, at t, I neither believe a nor $\neg a$, and neither perceive a nor $\neg a$. Then the second axiom gives both

$$\mathsf{B}\ (\neg a)(t+1)$$

and

$$\mathsf{B}\ (\neg\neg a)(t+1)$$

which is not an inconsistency in the logic but describes an inconsistent state of belief.

A fundamental part of our theory is that perceptibles are atomic, so α is used to range over atomic sentences only. This is clear in the paper (e.g. the semantics of the perception operator), but clearly wasn't in my (necessarily) hurried presentation.

This axiom might be considered unintuitive for the same reason as the negative introspection for belief is considered unintuitive. I don't perceive a rhinocerous in the room, so by default I believe that there is not a rhinocerous in the room. This seems an odd conclusion, as I was not aware/thinking of rhinoceri, this concept was not "cognitively active". In the case of belief, Halpern and (I believe Vardi) introduce an awareness operator and use this to distinguish between explicit belief and implicit belief. We can do a similar thing here to make the axiom more plausible. In our case, awareness would, as suggested, be something like "cognitively active". Thus I may not explicilty believe that there is a rhinocerous (Baby Mercedes, flying saucer, ...) in the room, because I am not currently using these concepts. However, I have the appropriate implicit beliefs, e.g. if asked whether there was a rhinocerous in the room, I would say that there wasn't one (and think that the question was a bit eccentric :-))

Any references on what, for want of a more precise term, I have called "cognitively active", would be appreciated.

Q2. **Pat Hayes**

You are proposing a formal common sense theory of perception and belief. However, in your paper there is very little about perception, and what you say is plain wrong. At the same time, there is an enormous wealth of knowledge about perception in the psychological literature, which you do not make any use of. Why should we pay any attention to this paper at all? (Additional debate).

A2. **John Bell**

(Full answer not recorded; the answer used a distinction between what the robot sees and what it says it sees).

C2-1. **Murray Shanahan**

I want to defend Bell's approach: for the purpose of building robot, I'm not interested in the psychological aspects of perception or in how it's done in people. It is important, however, to understand the relationship between an agent's sensory data interpretation and its beliefs, and that seems to be what the paper is about.

Q3. **Erik Sandewall**

You claim to present "a formal common sense theory of perception and belief" (first sentence of the abstract), but in the paper you suggest "that we perceive perceptual attributes". Do you also claim that all major aspects of perception can be characterized in terms of perceptual attributes? Otherwise, can you say something about what aspects of perception are covered and what are not?

(Answer not recorded).

Vol. 2: *105-106*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/330/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

John Bell

Chronological Minimization and Explanation

Q1. Workshop discussion

There were a number of comments by several participants on the question of chronological minimization vs chronological ignorance, and the author's reference to Shoham's work. The bottom line seemed to be that we have to make distinction between *chronological minimization of change* which was proposed by several authors concurrently in the flurry of the Yale shooting problem, and *chronological minimization of ignorance* which was developed by Yoav Shoham in his thesis and subsequent book. The early formulations of chron.min. of change were by Shoham (conference article), Kautz, and Lifschitz.

Several authors have investigated chron.min. of change in more recent work, whereas chron.min. of ignorance has not been further used. There are also indications that the latter gives incorrect results for many simple examples. However, the present article does seem to be using chron.min. of ignorance.

Q2. Erik Sandewall

I have objections to some parts of your argumentation at the beginning of the article, and finally some questions.

In the article, you write that I have "attacked" Shoham's causal theories for technical reasons, and that I "went on to propose increasingly more complex forms of chronological minimization", with reference to my 1994 book. This is inaccurate, since my approach in [mb-Sandewall-94] as well as elsewhere has been to compare and analyse different entailment methods, rather than advocating one of them over another. In particular, the book contains assessments for about one dozen such methods, several of which have nothing to do with chronological minimization. The set includes both a number of methods that had previously been introduced by others, and methods that naturally arose in the course of the work.

In the article, you also object to the use of filtering, where observations are imposed after the minimization of changes (or of discontinuities) has occurred. This method is now fairly widely used; it was adopted by Lifschitz under the name of nested circumscription. Your objection is that it is complex and that it does not have any counterpart in everyday causal reasoning. However, if simplicity is an issue, then I can offer it through PMON, also defined in [mb-Sandewall-94], which is about as simple as you can get.

With respect to having a counterpart in everyday causal reasoning, this introduces an issue that the book makes no attempt and no claim to address. The criticism is therefore irrelevant. The analyses of chron.min. as well as of PMON in the book only consider their use for the plain inertia problem; it is nowhere claimed that the analysis would apply for any kind of causal reasoning.

I am afraid that at this point you share in a habit that is quite common in our field, namely to criticize an approach X only because it fails to handle some example correctly, and regardless of whether it was *claimed* or *intended* to handle that case. In other words, one can be criticized for not solving all the problems at once. This style of discussion may be understandable in those cases where the author has not specified what class of problems his approach *is* claimed to solve in the first place, of if the author has made excessive claims. However, since I have tried to make a point of specifying what topics I am addressing at each point, I find your objection misdirected.

With respect to "everyday reasoning", please see my position statement (forthcoming) in the panel discussion.

These were my objections to your argumentation. Finally, a few questions:

1. In view of your reference to "everyday causal reasoning" in the context of my work, is the logic proposed in your article intended to match such reasoning with respect to both the structure of the inference process and the conclusions obtained? In particular, what sense does it then make to treat prediction and explanation as two different reasoning processes, as you propose to? After all, in everyday reasoning we are able to combine the two seamlessly.

2. In your article, you introduce the predicates *Aff* and *Qual* to be used in representing inertia and change. How are these related to similar predicates previously proposed by others? The *Aff* predicate seems to be occlusion, once again reinvented, and your general inertia rule seems to be the same as the nochange axiom schema used in PMON (originally introduced by McCarthy I believe).

3. The most important difference between your approach and what has been proposed before appears to be that you are using a threevalued logic; is this correct? If so, and in view of the similarity of other aspects of the approach to previous work, what advantage do you derive from switching to a three-valued logic for this purpose?

References:

[mb-Sandewall-94] Erik Sandewall. *Features and Fluents. The Representation of Knowledge about Dynamical Systems.* Oxford University Press, 1994.

Electronic News Journal on Reasoning about Actions and Change 107
Vol. 2: *107-109*, 28.2.1998
http://www.ida.liu.se/ext/etai/nj/fcs-98/349/

Protocol of the discussion at the FCS-98 workshop and of the subsequent on-line discussion during February, 1998 about the following research article:

Javier Pinto

Causality in Theories of Action

Q1. Graham White

Assume that a child is able to deal with switches and lights. Do you believe that the child has to have any knowledge of electrical current for this to be possible?

A1. Javier Pinto (29.1)

My work is not aimed towards dealing with children cognitive processes, therefore my answer is not based on any experimental evidence; rather, it is based on my own introspection and on my little experience with children development. I don't think that anybody could correctly analyze the behavior of the circuit without having some notion of what an electrical current is. Explanations based on other hypotheses might also be possible. The point of the paper is that if your explanations are based on explicit causal information, and if you do not get the right answers, you might want to change your model of the world to better reflect its behavior. Therefore, if a child gets the behavior wrong, because of a wrong analysis of the causal mechanism, the child should revise its knowledge, as opposed to revise his underlying reasoning mechanism to deal with causation.

Q2. Erik Sandewall

The representation examples that you have given are the classical ones, which others have adequate solutions for as well. Are there some examples which your representation does correctly and which others can not do right?

Also, the examples in the article are very simple ones. What is the most complex example that you have done in this representation?

A2. Javier Pinto (29.1)

The most complex examples are the standard examples that appear in the literature related to causality (along with some examples involving continuous change, common in the qualitative reasoning literature). I believe that the work of other people can also be seen as instances of what I propose (e.g. the work of Todd Kelley on qualitative physics). In particular, we have looked at extensions to the circuit example (as illustrated in the paper and taken from Thielscher [j-aij-89-317]. Also, as shown in the article, we use a similar approach to deal with indeterminacy. I admit that this latter treatment of indeterminacy might be regarded as too complex. This added complexity is unavoidable, since we need to worry about the many possible outcomes of an indeterminate event, and must also be prepared to reason about each possible outcome. In future work I plan to study whether one can show a formal correspondence between approaches based on explicit causation and my own. My conjecture is that there is always a translation from an explicit causal representation and the one i propose. However, it remains to be seen if this translation is a natural one.

References:

[j-aij-89-317] Michael Thielscher. *Ramification and causality.* Artificial Intelligence, vol. 89 (1997), pp. 317-364.

Q3. **Tom Costello**

People are able to verify systems as complex as the Pentium, why then worry about these silly little circuits?

A3. **Javier Pinto** (29.1)

Well, I guess the best answer to this was given by John McCarthy, i.e., these sort of examples are the *drosophilas* of AI and Knowledge Representation. That is, the examples represent problems that are good testbeds for the different essential problems that we confront in AI.

Thus, these simple examples allow us to isolate a particular problem that we want to deal with. We do not need a zillion transistors to model the problems that arise when dealing with ramifications. In particular, this problem illustrates reasoning problems related to *causality*. There are many variants of this example, like the *stuffy room* example of Ginsberg and Smith [reasacti] and the suitcase with two latches of Lin [c-ijcai-95-1985]. All these examples have allowed us to illustrate problems with our reasoning based on formalizations of these domains, and have helped us advance the state of the art in the area.

References:

[c-ijcai-95-1985] Fangzhen Lin. *Embracing Causality in Specifying the Indirect Effects of Actions.* Proc. International Joint Conference on Artificial Intelligence, 1995, pp. 1985-1991.

C3-1. **Tom Costello**

What if the switches are very close, or if ...?

C3-2. **Javier Pinto** (29.1)

The point of the question, as I understand it is that the model proposed did not account with the many different ways in which the model might be a mistaken idealization of reality. Certainly the model proposed is not geared towards dealing that particular issue. Rather, assuming that the circuit behaves the way it does, and that the primitive actions are independent flippings of the circuit, then how can we reason about it? Therefore, the answer is that we cannot deal with those situations since we were trying to deal with a completely separate set of issues.