Linköping Electronic Newsletters

ECSTER Newsletter

November-December, 1996

Erik Sandewall (ed.)

Linköping University
Linköping, Sweden

http://vir.liu.se/brs/news/96/dec/

Issued on 23.12.1996 by
Department of Computer and Information Science
Linköping University
581 83 Linköping, Sweden

Editor: Erik Sandewall

Typeset using LaTeX
Formatted using étendu *style*

It is intended that the present issue may be
published electronically in a persistent fashion
using Linköping University Electronic Press.

The URL stated at the bottom of the first page
will persistently contain information about the
publication status of the present newsletter issue.
The general URL for the ECSTER Electronic Colloquium is
`http://vir.liu.se/brs/`

This issue of the ECSTER Newsletter features the following news items:

1. **Under the auspices of ECCAI, a proposal has been developed for a new, European publication system for Artificial Intelligence.**
2. **ECSTER starts a colloquium debate - positions statements by Shanahan, Miller, and Lifschitz in this issue. Contributions invited from all readers!**
3. **Lifschitz starts "Page of Positive Reviews"**
4. **Linköping University Electronic Press has started operation**

These news items have one thing in common: they are signs of an on-going shift in publication patterns, due to the existence of the Internet as a viable publication mechanism. It may lead to much bigger changes than merely the use of electronic transmission instead of paper and post office for sending the articles from author to reader; it may also lead us to reconsider the present peer-review mechanism which precedes publication.

The familiar publication pattern for research results, based on anonymous peer reviw prior to publication in an established journal or - particularly in computer science - in a quality conference proceedings, has dominated during the last 50 years. A number of questions are now being asked by researchers in many disciplines as online publication emerges and finds its own best modus operandi. Do we need printed journals? Do we need publishers? How shall we deal with the long publication delays? Should peer review be done before or after first publication? What should be the balance between open and anonymous review?

The recent developments which are reported in the present issue of the ECSTER Newsletter have this in common: they all indicate new directions that are made possible by online, Internet-based exchange of research results.

1 Proposal for a new, European publication scheme

The recent ECAI (European Conference on AI) dedicated an evening session to the question of a possible new, European AI journal which would be organized in a novel way and make the best use of the electronic medium. Proceeding from the recommendations that were made during that evening session and the discussions that have followed it, a concrete proposal has now been developed. It has been written by Erik Sandewall with discussions, in particular, with Wolfgang Wahlster and Georg Gottlob.

The basic idea in this proposal is to distinguish between a paper-based journal - the AICom - and an electronic medium: the Electronic Transactions on Artificial Intelligence, or **ETAI**. The ETAI will be based on a principle of **open, posteriori review**: articles are *first published* based only on local self-control within each research group; then they are subjected to *open review* (non-anonymous) using electronic discussion groups, and after a certain period of open review they can be considered for *certification*, that is, for promotion to a status corresponding to conventional journal publication today.

The ECSTER Newsletter and, more broadly, the services of the ECSTER colloquium have served as an example of what can be done with the electronic medium. It is observed that similar structures have emerged in other branches of AI as well, and it is foreseen that such activities will continue to expand.

The proposed ETAI publication scheme offers a number of advantages for the authors. By considering research results as published prior to review, it helps to protect the author's priority right to his or her results. By allowing for open review, it facilitates for authors of earlier, related work to contribute their comments, critical as well as constructive, and thus to influence the article. By separating publication from certification, it makes it possible for reviewing mistakes to be corrected: even if the certification of an article is delayed, the author still retains his or her original date of publication when the work is finally recognized.

The [1][full proposal] is of course available online. There is also an [2][accompanying article] discussing the general publishing principles which underly the proposal.

2 Page for positive, open reviews started by Lifschitz

In a move which parallels the ETAI initiative, Vladimir Lifschitz has started a [3][page of positive reviews] (PPR) which features open reviewing of research articles in logicist AI. At present (December 16), the PPR has collected 8 reviews. Although these reviews are in fact *very* positive, one can easily see a possibility that reviews containing critique as well as praise could be published on the PPR, and that interesting debates could emerge. The present editor considers the PPR as an indication of the right way to go.

[1]The URL is http://www.ida.liu.se/ erisa/org/spai/ntd96/05/
[2]The URL is http://www.ida.liu.se/ erisa/org/spai/ntd96/04/
[3]The URL is http://ppr.kr.org/ppr/

3 ECSTER colloquium debate on approaches to reasoning about actions and change

Several approaches to reasoning about actions and change co-exist at present in the literature. The major divide seems to be between the situation calculus on one hand, and approaches using explicit time on the other hand. It may not be easy for the readers of this literature to see how the different approaches relate, and what are their respective weaknesses and strengths. Sometimes, it is even difficult for the researchers in the area to make this analysis. For example, in this recent KR paper, Ray Reiter writes:

There have been a few earlier papers on formalizing natural actions and continuous time. Shanahan's approach [30] is embedded in the **event calculus** *(Kowalski and Sergot [11]); Sandewall [27] relies on a temporal logic. Accordingly, these proposals are difficult to compare with ours, based as it is on the situation calculus.*

After a suggestion by Ray, ECSTER invites researchers in this area to an on-line colloquium exchange of views on **different approaches to reasoning about actions and change**. The purpose of the debate is *to clarify what are the major alternative approaches to reasoning about actions and change in contemporary research*, and also *to identify and compare the capabilities and the limitations of those approaches*.

A separate [4][debate page] has been set up and will contain successive debate contributions. It presently contains an introduction by Erik Sandewall and position statements by Murray Shanahan, Rob Miller, and Vladimir Lifschitz. All contributions which are sent to the present editor will be added to the debate page.

Also, for readers who wish to receive each debate contribution as an E-mail message, we are going to set up a mailgroup. Send a message to the [5][Newsletter editor] in order to be included in this mailgroup.

(Note: the service will be closed between December 28 and January 2 due to vacation travel).

[4]The URL is http://vir.liu.se/brs/news/96deb/contents.html

[5]E-mail address: erisa@ida.liu.se

4 Initial statements in the Colloquium Debate

4.1 Current Research on Reasoning About Actions and Change: Topics for a Debate

4.1.1 Erik Sandewall

Several approaches to reasoning about actions and change co-exist at present in the literature. The major divide seems to be between the situation calculus on one hand, and approaches using explicit time on the other hand. It may not be easy for the readers of this literature to see how the different approaches relate, and what are their respective weaknesses and strengths. Sometimes, it is even difficult for the researchers in the area to make this analysis. For example, in this recent KR paper, Ray Reiter writes: [Continued on a later page in this document]

4.2 Reasoning about Actions: A Position Statement

4.2.1 Murray Shanahan

AI needs an action formalism that is expressive, and that incorporates a solution to the frame problem that's robust in the face of the phenomena it can represent. The formalism should be expressive enough to represent at least the following phenomena.

1. Actions with indirect effects (ramifications)
2. Concurrent action
3. Non-deterministic action
4. Narrative time
5. Continuous change

A rigorous argument that the formalism in question solves the frame problem should be supplied.

Here comes the controversial bit. [Continued on a later page in this document]

4.3 Comparing Action Formalisms: A Position Statement

4.3.1 Rob Miller

Here are some fairly miscellaneous thoughts about comparing alternative approaches to Reasoning about Action. ...

(1) When comparing and evaluating formalisms, we need to be careful not to form too strong associations between particular *methodologies* (e.g. deduction and entailment methods, default reasoning techniques) and particular *ontologies*. I can think of a few occasions...

(2) As a community, we should be encouraging work on comparing action formalisms and ontologies, and we should be critical of papers which don't contain adequate comparisons with other work (and especially with work based on different ontologies). There is now a fair body of work exploring how the Event Calculus and the Situation Calculus correspond, so there's really no excuse for lack of comparisons in this case at least. [Continued on a later page in this document]

4.4 Approaches to Reasoning About Actions: A Position Statement

4.4.1 Vladimir Lifschitz

1. Explicit time vs. the situation calculus. The following situation calculus formula seems to have no counterpart in languages with explicit time:

value(f,result(a1,s)) = value(f,result(a2,s)). (1)

It says that the value of f at the next instant of time does not depend on which of the actions *a1, a2* is going to be executed. For instance,...

4. Why are there so many action languages? An action language is a formal model of the part of natural language that is used for describing the effects of actions. Whenever we improve our understanding of that part of natural language, this improved understanding may be expressed by defining a new dialect of "script-A." I expect that...

5. Explicit information about causal directions. Causality differs from material implication in that it is not contrapositive...

[Continued on a later page in this document]

5 News from the Linköping scene

5.1 Linköping University Electronic Press starts operation

Linköping University has set up a special organization for electronic on-line publication, called the [6][Linköping University Electronic Press], or **E-Press** for short. Its primary mission is to perform *unrefereed electronic publishing* of scientific articles - like a preprint archive, with a guarantee that what has been published there remains available.

The existence of organizations like the Linköping E-Press is presumed by the proposed publication scheme for ETAI. One important part of the idea is that they can make research articles available electronically at the expense of the *authoring institution*, and without a need to charge the reader for the access.

5.2 Major grant awarded by the Wallenberg Foundation

The Knut and Alice Wallenberg Foundation has awarded a grant of 40 million Sw.Crowns (about 5 million ECU) for research on [7][**Information Technology for Autonomous Aircraft**] during an initial period of three years. The project will be coordinated by Erik Sandewall, and will consist of four sub-projects; the sub-project for **high-level autonomous decisions** will include reasoning about actions, spatial, and temporal reasoning, and is therefore directly related to ECSTER interests. Other subprojects address computer vision, VLSI design, and simulation.

6 Articles first published during November-December, 1996

A section with this kind of heading will be used for advertising articles which have been made available through first publication entities, preprint archives, etc. since the previous issue of the Newsletter. This is in the spirit of the proposal for the ETAI research publication scheme which was described above.

- Erik Sandewall.
  Underlying Semantics for Action and Change with Ramification. Published by Linköping University Electronic Press, `http://www.ep.liu.se/ea/cis/1996/002/`. Accepted for inclusion in: Oliviero Stock (ed): Spatial and Temporal Reasoning. Kluwer Academic Publishers, 1997.

- Erik Sandewall.
  Towards the validation of high-level action descriptions

[6] The URL is http://www.ep.liu.se/
[7] The URL is http://www.ida.liu.se/ext/witas/index-eng.html

from their low-level definitions. Published by Linköping University Electronic Press, `http://www.ep.liu.se/ea/cis/1996/004/`. Accepted for inclusion in AICom, December 1996 issue, published by IOP Press.

Current Research on Reasoning about Actions and Change: Topics for a Debate

Erik Sandewall
Linköping University, Sweden

Several approaches to reasoning about actions and change co-exist at present in the literature. The major divide seems to be between the situation calculus on one hand, and approaches using explicit time on the other hand. It may not be easy for the readers of this literature to see how the different approaches relate, and what are their respective weaknesses and strengths. Sometimes, it is even difficult for the researchers in the area to make this analysis. For example, in this recent KR paper, Ray Reiter writes:

There have been a few earlier papers on formalizing natural actions and continuous time. Shanahan's approach [30] is embedded in the **event calculus** *(Kowalski and Sergot [11]); Sandewall [27] relies on a temporal logic. Accordingly, these proposals are difficult to compare with ours, based as it is on the situation calculus.*

After a suggestion by Ray, ECSTER invites researches in this area to an on-line colloquium exchange of views on **different approaches to reasoning about actions and change**. The basic idea is to have a mailgroup which is combined with a lasting on-line presentation of the accumulated contributions, and with a permanent publication of the entire debate.

The purpose of the debate is *to clarify what are the major alternative approaches to reasoning about actions and change* in contemporary research, and also *to identify and compare the capabilities and the limitations of those approaches*.

Some distinctions will be made already at this point in order to further define the topic. We propose a distinction on ontological grounds between *situation calculus approaches* and *narrative timeline approaches*, which are defined as follows: In narrative timeline approaches, one uses a multisorted logic where "time" is one of the sorts, and actions are *attached to* the timeline using a construct such as *Do(s,t,a)*, saying that the action *a* is performed during the interval starting at *s* and ending at *t*. In situation calculus approaches, on the other hand, one uses an equally multisorted logic where "situations", as one of the sorts, form a tree-structured domain where each situation *contains* a sequence or other structure of actions. Thus, to express that the property *p* holds when the action *a* is concluded, a situation-calculus approach would write something of the form

Holds(p, Result(a,s)), and a narrative timeline approach would write something along the lines of *Do(s,t) and Holds(p,t)*. In both cases, there are of course many variants to the theme. One topic for the present debate is what are the advantages and disadvantages of these two approaches.

Within each of those approaches, and possibly independently of the distinction, there are various *entailment methods* which define how to obtain the intended conclusions for a given scenario description. Some of these entailment methods are defined in terms of preference relations or other selection mechanisms on models, others are defined in terms of syntactic transformations on the initially given set of axioms. Chronological minimization of change is an example of a semantically defined method; explanation closure is an exampleof a method defined through syntactic transformations. One topic for the present debate is what are the presently available entailment methods (including both those that are defined semantically and those defined syntactically) and what is known about their properties.

A number of techniques which have been proposed in recent years have been adopted by several researchers. These techniques include the use of *occlusion*, *filtering*, *nested circumscription*, the *release predicate*, and *composition of actions*, but the list can probably be extended. One topic for the present debate is what are these generally used techniques, and to identify cases where a previously known technique reappears in new guise or disguise.

The concepts of *intended models*, and of an *underlying semantics* defining the set of intended models, have developed as a way of characterizing what one expects from a logic of actions and change. This raises a number of topics for the present debate: what are appropriate ways of defining intended models; in what sense are intended models truly "intended"; are there alternative definitions of intended models and how do they relate to each other.

One noticable phenomenon in recent years has been the appearance of *action description languages*, in particular the different variants of the script-A language. Some questions of debate are: in what ways are action description languages different from logics (or are they?); why are there so many action description languages; and how do action description languages relate to underlying semantics.

When *ramification* is addressed, there is an issue between those methods using minimization of change and those methods that make use of explicit information about causal directions. What is true about the capabilities and limitations of these alternatives?

This is already a number of non-trivial questions, but that should not preclude anyone from also addressing other questions of a similar character with respect to reasoning about actions and change.

Reasoning About Actions: A Position Statement

Murray Shanahan
Queen Mary and Westfield College, London, England

Since this is a position statement, I suppose it's legitimate to give a sales pitch for a particular formalism, and to be a bit controversial.

AI needs an action formalism that is expressive, and that incorporates a solution to the frame problem that's robust in the face of the phenomena it can represent. The formalism should be expressive enough to represent at least the following phenomena.

1. Actions with indirect effects (ramifications)
2. Concurrent action
3. Non-deterministic action
4. Narrative time
5. Continuous change

A rigorous argument that the formalism in question solves the frame problem should be supplied.

Here comes the controversial bit. I believe that the most instructive way to validate a formalism is through a judiciously chosen set of representative benchmark scenarios. I sometimes feel that attempts to do this validation by proving a relationship between a formalism and some other formal structure are just an excuse for introducing a lot of unnecessary mathematics. I strongly believe that we can only contribute to AI if we divert our efforts away from proving "interesting" theorems and into the application of our formalisms to fundamental problems in AI (such as planning in robots).

Here comes the sales pitch. The event calculus presented in my forthcoming book is capable of representing all the above mentioned phenomena, and the solution to the frame problem that accompanies it is provably immune to the Hanks-McDermott problem. In recent ECAI and AAAI papers, this formalism is applied to a serious example involving a robot. The example includes all five of the above listed representational features. If your favourite action formalism can't do this robot example by now, maybe it's time to change to a new one.

Honesty time now. The event calculus cannot, as yet, handle knowledge producing actions, or complex actions (ie: actions including program constructs). To me these requirements seemed less urgent than the ones in my list, but they're undoubtedly important.

Finally, the last thing we need is new formalisms. The only excuse for introducing a new formalism is that it is fundamentally different from any of the existing ones. Instead, we have to try to expose the underlying similarities and differences between possible formalisms in the hope that we can start to understand the range of representational choices in a principled way.

The above opinions are those of the author on 20th November 1996, and may be subjected to total revision when he is older and wiser, or possibly sooner.

Comparing Action Formalisms: A Preliminary Position Statement

Rob Miller
Imperial College, London, England

Here are some fairly miscellaneous thoughts about comparing alternative approaches to Reasoning about Action. (Please forgive me for a few fairly blatant attempts to advertise my own work at the same time.)

(1) When comparing and evaluating formalisms, we need to be careful not to form too strong associations between particular *methodologies* (e.g. deduction and entailment methods, default reasoning techniques) and particular *ontologies*. I can think of a few occasions when this has happened. For example: (i) At least until Murray Shanahan's 1995 paper [8], the ontology of the Event Calculus was irrevocably linked in many researchers' minds with Logic Programming (and in particular with negation-as-failure), and thus dismissed or attacked on "semantic" grounds. (ii) The methodology of using "action languages" (i.e. the Language A, etc. [1]) has become overly associated with the ontology that the original Language A inherited from the Situation Calculus. (Tony Kakas's and my "Language E" paper [2] has shown that, for better or for worse, the methodology can actually be applied in a wider context.) (iii) It seems to be a commonly held belief that "narrative time" ontologies, such as that of the Event Calculus, demand that planning be done using abduction (as opposed to deduction). (I've shown that this is not the case in [6].)

(2) As a community, we should be encouraging work on comparing action formalisms and ontologies, and we should be critical of papers which don't contain adequate comparisons with other work (and especially with work based on different ontologies). There is now a fair body of work exploring how the Event Calculus and the Situation Calculus correspond, so there's really no excuse for lack of comparisons in this case at least. (For formal results, see for example [3], [5], [7], [9] and [2] - the last shows a correspondence between Languages A and E.)

(3) In [2], Tony Kakas and I wrote:

"We believe that the use of, and comparison between, different ontolgies is vital in the study of reasoning about action. Central issues such as the frame problem, the ramification problem and the qualification problem all take on different flavours when set in different onto-

logical contexts. Comparisons between different approaches can help reveal which aspects of these problems are fundamental, and which are merely the product of a particular method of representation".

I stand by this view. A good example of a (nevertheless interesting) problem which is the product of a particular ontology (rather than being fundamental) is the difficulty of distinguishing between observations and causal rules in the Situation Calculus and in the Language A (i.e. in the context of the Language A, the difficulty in distinguishing the roles of value and effect propositions). Vladimir Lifschitz presented a technically interesting solution to this difficulty in [4], but neither the problem nor the solution translate to other (ontologically different) approaches. We need to be careful to distinguish between this type of issue and more fundamental problems such as dealing with ramifications or continuous change.

(4) We need to keep the role of "action languages" (the Language A, etc.) in perspective. To quote Vladimir Lifschitz [4]:

"Originally, action languages were meant to play an auxiliary role. The primary goal was to represent properties of actions in less specialised formalisms, such as first-order logic and its nonmonotonic extensions, and the idea was to present methods for doing that as translations from action languages".

Well, we shouldn't loose sight of that primary goal. There are many good reasons for using a general purpose logic to represent properties of actions. Perhaps the most important is that it allows us to link in with work on other aspects of common sense reasoning (reasoning about space, shape, beliefs, contexts, etc.).

REFERENCES

[1] Michael Gelfond and Vladimir Lifschitz, *Representing Actions in Extended Logic Programming*, Proceedings of the Joint International Conference and Symposium on Logic Programming, ed. Krzysztof Apt, MIT Press, page 560, 1992.

[2] Antonios Kakas and Rob Miller, *A Simple Declarative Language for Describing Narratives with Actions*, to appear in the Journal of Logic Programming: Special Issue on Reasoning about Action and Change (scheduled for publication by Elsevier Science early 1997). [8][[Abstract, Postscript and DVI Versions]]

[3] Robert A. Kowalski and Fariba Sadri, *The Situation Calculus and Event Calculus Compared*, in Proceedings of the International Logic Programming Symposium (ILPS'94), 1994.

[4] Vladimir Lifschitz, *Two Components of an Action Language*, in Working Papers of the 3rd Symposium on Logical Formalizations of Commonsense Reasoning (Common Sense '96), Stanford University, 1996.

[5] Rob Miller, *Situation Calculus Specifications for Event Calculus Logic Programs*, in Proceedings of the Third International Conference on Logic Programming and Non-monotonic Reasoning, Lex-

[8]The URL is http://www-lp.doc.ic.ac.uk/UserPages/staff/rsm/abstract9.html

ington, KY, USA, pub. Springer Verlag, 1995, pages 217-230. [9][[Abstract, Postscript and DVI Versions]]

[6] Rob Miller, *Notes on Deductive and Abductive Planning in the Event Calculus*, July, 1996. [10][[Abstract, Postscript and DVI Versions]].

[7] Alessandro Provetti, *Hypothetical Reasoning about Actions: From Situation Calculus to Event Calculus*, Computational Intelligence, volume 12, number 2, 1995.

[8] Murray Shanahan, *A Circumscriptive Calculus of Events*, Artificial Intelligence, vol. 77, pages 249-284, Elsevier Science Publishers, 1995.

[9] Kristof Van Belleghem, Marc Denecker and Danny De Schreye, *On the Relation Between Situation Calculus and Event Calculus*, to appear in the Journal of Logic Programming: Special Issue on Reasoning about Action and Change (scheduled for publication by Elsevier Science early 1997).

[9]The URL is http://www-lp.doc.ic.ac.uk/UserPages/staff/rsm/abstract3.html

[10]The URL is http://www-lp.doc.ic.ac.uk/UserPages/staff/rsm/abstract14.html

Approaches to Reasoning About Actions: A Position Statement

Vladimir Lifschitz
University of Texas at Austin, TX, USA

1. Explicit time vs. the situation calculus. The following situation calculus formula seems to have no counterpart in languages with explicit time:

value(f,result(a1,s)) = value(f,result(a2,s)). (1)

It says that the value of f at the next instant of time does not depend on which of the actions *a1, a2* is going to be executed. For instance, if I now send an e-mail message to Erik Sandewall, the total number of messages sent by me since this morning will be the same as if I send a message to Ray Reiter instead. This is an argument in favor of the situation calculus.

But there is a little problem here. What is the meaning of (1) if the effects of *a1* and *a2* on *f* are nondeterministic? I have a few coins in my pockets; let *a1* stand for getting a coin from my left pocket, let *a2* stand for getting a coin from my right pocket, and let *f* stand for the value of the coin that I have in my hand. We can interpret (1) as a counterfactual, but this seems less interesting than assertions involving some kind of quantification over the outcomes of *a1* and *a2*, for instance:

(i) there exist an outcome of *a1* and an outcome of *a2* such that (1) holds,

(ii) for any outcome of *a1* and any outcome of *a2*, (1) holds,

(iii) for any outcome of *a1* there exists an outcome of *a2* such that (1) holds.

The situation calculus has no mechanism for expressing these distinctions.

2. Filtering. I understand it as applying a nonmonotonic logic to a subset of the given facts and then using the remaining facts–"constraints"–to discard some of the models of this nonmonotonic theory. This is a powerful idea, and "nested abnormality theories" are merely a syntactic device that can be used to describe filtering. In reasoning about action, treating initial conditions as constraints makes the formalization problem easier. The reduction of domain circumscription to predicate circumscription in John McCarthy's 1980 paper is an early example of filtering.

3. Occlusion. I undertand it as restricting inertia so that it would not apply to some fluents at some instants of time. This is a

special case of the more general idea of restricting a default so that it would not apply to some objects. According to an axiom from John McCarthy's 1986 paper on applications of circumscription, unless an object is abnormal in aspect 1, it can't fly; then birds are declared to be "possible exceptions" to this default by postulating that they are abnormal in aspect 1. This is similar to occlusion.

4. Why are there so many action languages? An action language is a formal model of the part of natural language that is used for describing the effects of actions. Whenever we improve our understanding of that part of natural language, this improved understanding may be expressed by defining a new dialect of "script-A." I expect that we will see many such dialects in the future.

And I hope that this will bring us simplicity and elegance, rather than the multiplication of independent language constructs. In the first action language that was capable of representing ramifications, there were "causes" propositions for representing dynamic causal laws, and "always" propositions for representing static causal laws. But recent research on the logic of causality shows that causal laws of both kinds have similar properies; we can expect that in future action languages "causes" and "always" will be subsumed by a more general construct.

5. Explicit information about causal directions. Causality differs from material implication in that it is not contrapositive. The crucial role of this difference for the study of action and change is among the most important things that we have learned in this area over the last years. Several recent theories of causality grasp this distinction, but they do not seem to be mathematically reducible to each other. This is very much like what happened around 1980, when several mathematically non-equivalent nonmonotonic formalisms were proposed. Attempts to relate these formalisms to each other have led to interesting research in logic.