A Probabilistic Event Calculus

Ben Goertzel
Novamente LLC
November, 2006

1 Introduction

The Event Calculus (Kowalski and Sergo, 1986; Miller and Shanahan, 1991), a descendant of Situation Calculus (McCarthy and Hayes, 1969), is perhaps the best-fleshed-out attempt to apply predicate logic to the task of reasoning about commonsense events. A recent book by Erik Mueller (2006) reviews the application of Event Calculus to the solution of a variety of "commonsense inference problems," defined as simplified abstractions of real-world situations.

This brief, informal article describes a variation of Event Calculus called Probabilistic Event Calclus, in which the strict implications from standard Event Calculus are replaced with probabilistic implications, and other changes are also introduced, such as a repositioning of events and actions in the basic event ontology, and the introduction of a simpler mechanism to avoid the use of circumscription for avoiding the "frame problem." These changes make it much easier to use Event Calculus type reasoning within the Probabilistic Logic Networks (PLN) inference framework (Ikle' et al, 2006) created for use within the integrative Novamente AI Engine artificial general intelligence architecture (Goertzel, 2006).

More qualitatively, it is suggested that these changes also result in a more "cognitively natural" sort of Event Calculus; though of course this sort of claim is hard to substantiate rigorously, and we will not pursue this line of argument here, except to a few simple points:

- There is growing evidence for probabilistic calculations in the human brain, whereas neural bases for crisp predicate logic and higher-order logic mechanisms like circumscription have never been identified even preliminarily
- The use of probabilistic implications makes clearer how reasoning about events may interoperate with perceptions of events (given that perception is generally uncertain) and data mining of regularities from streams of perceptions of events (which also will generally produce regularities with uncertain truth values)
- As will be discussed in the final section, the pragmatic resolution of the "frame problem" seems more straightforward using a probabilistic variant of Event Calculus, in which different events can have different levels of persistence

2 A Simple Event Ontology

Probabilistic Event Calculus, as we define it here, involves the following categories of entity:

- events
  - fluents
  - general events
- temporal predicates
  - holding
  - initiation
  - termination
  - persistence
  - etc.

- actions
- time distributions
    - time points
    - time intervals
    - general time distributions

A "time distribution" refers to a probability density over the time axis; i.e., it assigns a probability value to each interval of time. Time points are considered pragmatically as time distributions that are bump-shaped and supported on a small interval around their mean (true instantaneity being an unrealistic notion both psychologically and physically). Time intervals are considered as time distributions corresponding to characteristic functions of intervals.

The probabilistic predicates utilized begin with the functions:

```
hold (event)
initiate(event)
terminate(event)
```

These functions are assumed to map from events into probabilistic predicates whose inputs are time distributions, and whose outputs are probabilistic truth values. The class of "events" may be considered pragmatically as the domain of these functions.

Based on these three basic functions, we may construct probabilistic predicates such as

```
holdsAt(event, time point) = (hold (event))(time point)
initiatedAt(event,time point) = (initiate(event))(time point)
terminatedAt(event, time point) = (terminate(event))(time point)

holdsThroughout(event, time interval) = (hold (event))(time interval)
initiatedThroughout(event, time interval) = (initiate (event))(time interval)
terminatedThroughout(event, time interval) = (terminate (event))(time interval)

holdsSometimeIn(event, time interval)  =
    "There exists a time point t in time interval T so that holdsAt(E,t)"
initiatedSometimeIn(event, time interval)
    "There exists a time point t in time interval T so that initiatedAt(E,t)"
terminatedSometimeIn(event, time interval)
    "There exists a time point t in time interval T so that terminatedAt(E,t)"
```

It may seem at first that the interval-based predicates could all be defined in terms of the time-point-based predicates using universal and existential quantification, but this isn't quite the case. Initiation and termination may sometimes be considered as processes occupying non-instantaneous stretches of time; so that a process initiating over an interval does not imply that process initiating at each point within that interval.

Next, there are various important properties that may be associated with events, for example persistence and continuity.

```
Persistent(event)
Continuous(event)
Increasing(event)
Decreasing(event)
```

which are (like hold, initiate and terminate) functions outputting probabilistic predicates mapping time distributions into probabilistic truth values.

Persistence indicates that the truth value of an event can be expected to remain roughly constant over time from the point at which is initiated until the point at which the event is terminated. . A "fluent" is then defined as an event that is persistent throughout its lifetime.

Continuous, Increasing and Decreasing apply to non-persistent events, and indicate that the truth value of the event can be expected to {vary continuously, increase, or decrease} over time.

For example, to say that the event of clutching (e.g. the agent clutching a ball) is persistent involves the predicate (isPersistent(clutching))([-infinity,infinity]). Note that this predicate may be persistent throughout all time even if it is not true throughout all time; the property of persistence just says that once the event is initiated its truth value remains roughly constant until it is terminated.

Other temporal predicates may be defined in terms of these. For example, an "action" may be defined as an initiation and termination of some event that is associated with some agent (which is different from the standard Event Calculus definition of action).

Next, there is also use for further derived constructs such as

```
initiates(action, event)
```

indicating that a certain action initiates a certain event or

```
done(event)
```

which is true at time-point t if the event terminated before t (i.e. before the support of the time distribution representing t).

Finally, it is worth noting that in a logic system like PLN, the above predicates may be nested within markers indicating them as hypothetical knowledge. This enables Probabilistic Event Calculus to be utilized much like Situation Calculus (McCarthy, 1986), in which hypothetical events play a critical role.

3 The Frame Problem, and the Commonsense Law of Inertia

A major problem in practical applications of Event Calculus is the "frame problem" (McCarthy, 1986; Mueller, 2006), which – as usually construed in AI -- refers to the problem of giving AI reasoning systems implicit knowledge about which aspects of a real-world situation should be assumed *not* to change during a certain period of time. More generally, in philosophy the "frame problem" may be construed as the problem of how a rational mind should bound the set of beliefs to change when a certain action is performed. This section contains some brief conceptual comments on the frame problem and its relationship to Probabilistic Event Calculus.

For instance, if I tell you that I am in a room with a table in the center of it and four chairs around it, and then one of the chairs falls down, you will naturally assume that the three other chairs did not also fall down – and also that, for instance, the whole house didn't fall down as well (perhaps because of an earthquake). There are really two points here:

1. The assumption that, unless there is some special reason to believe otherwise, objects will generally stay where they are; this is an aspect of what is sometimes known as the "commonsense law of inertia" (Mueller, 2006).
2. The fact that, even though the above assumption is often violated in reality, it is beneficial to assume it holds for the sake of making inference tractable. The inferential conclusions obtained may then be used, or not, in any particular case depending on whether the underlying assumptions apply there.

The original strategy John McCarthy proposed for solving the frame problem (at least partially) was to introduce the formal-logical notion of circumscription (McCarthy, 1986). For example, if we know

```
initiatesDuring(chair falls down, T1)
```

regarding some time interval T1, then the circumscription of holdsDuring and T1 in this formula is

```
initiatesDuring(x,T1) <==> x = "chair falls down"
```

Basically this just is a fancy mathematical way of saying that no other events are initiated in this interval, except the one event of the chair falling down. If there are multiple events initiated during the interval, then one can circumscribe the combination of events, arriving at the assertion that no other events byt the ones in the given set occur. This approach has known shortcomings which have been worked around via various mechanisms, including the simple addition of an axiom stating that an events are by default persistent in the sense given above (see Reiter, 1991; Sandewall, 1998). Mueller (2006) uses circumscription together with workarounds to avoid the problems classically found with it.

However, none of these mechanisms is really satisfactory. In a real-world scenario there are always various things happening; one can't simply assume that nothing else happens except a few key events one wants to reason about. Rather, a more pragmatic approach is to assume, for the purpose of doing an inference, that nothing important and unexpected happens that is directly relevant the relationships one is reasoning about. Event persistence must be assumed probabilistically rather than crisply; and just as critically, it needs be assumed only for appropriate properties of appropriate events that are known or suspected to be closely related to the events being reasoned about in a given act of reasoning.

This latter issue (the constraints on the assumption of persistence) is not touched upon in most treatments of formal commonsense reasoning, because these treatments handle "toy domains" in which reasoning engines are fed small numbers of axioms and asked to reason upon them. This is quite different than the situation of an embodied agent that receives a massive stream of data from its sensors at nearly all times, and must define its own reasoning problems and its own relevant contexts. Thus, the real trickiness of the "frame problem" is not exactly what the logical-AI community has generally made it out to be; they have sidestepped the main problem due to their focus on toy problems.

Once a relevant context is identified, it is relatively straightforward for an AI reasoning system to say "Let us, for the sake of drawing a relatively straightforward inference, make the provisional assumption that all events of appropriate category in this context (e.g. perhaps: all events involving spatial location of inanimate household objects) are persistent unless specified otherwise." Information about persistence doesn't have to be explicitly articulated about each relevant object in the context, any more than an AI system needs to explicitly record the knowledge that each human has legs -- it can derive that Ben has legs from the fact that most humans have legs; and it can derive that Ben's refrigerator is stationary from the fact that most household objects are stationary. The hard part is actually identifying the relevant context, and understanding the relevant categories (e.g. refrigerators don't move around much, but people do). This must be done inductively, e.g. by knowledge of what contexts have been useful for similar inferences in the past. This is the crux of the frame problem:

- understanding what sorts of properties of what sorts of objects tend to be persistent in what contexts (i.e., learning specific empirical probabilistic patterns regarding the Persistent predicate mentioned above)
- understanding what is a natural context to use for modeling persistence, in the context of a particular inference (e.g. if reasoning about what happens indoors, one can ignore the out of doors even it's just a few feet away through the wall, because interactions between the indoors and out of doors occur only infrequently)

And this, it seems, is just plain old "AI inference" – not necessarily easy, but without any obvious specialness related to the temporal nature of the content material.

References

- B. Goertzel (2006). Patterns, Hypergraphs and General Intelligence. Proceedings of IJCNN-2006, Workshop on Human-Level Intelligence.
- M. Ikle', Ben Goertzel and Izabela Goertzel (2006). Quantifying Weight of Evidence in Uncertain Inference via Hybridizing Confidence Intervals and Imprecise Probabilities, Proceedings of 2006 AGI Workshop, IOS Press, to appear
- R. Kowalski and M.Sergot (1986) "A Logic-Based Calculus of Events" New Generation Computing, vol. 4 pp. 67-95.
- J. McCarthy (1986). Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28:89-116.
- J. McCarthy and P. J. Hayes (1969). Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463-502.

- E. T. Mueller (2006). *Commonsense Reasoning*. San Francisco: Morgan Kaufmann.
- R. Miller and M. Shanahan (1999). The event-calculus in classical logic - alternative axiomatizations. *Electronic Transactions on Artificial Intelligence*, 3(1):77-105.
- R. Reiter (1991). The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression. In Vladimir Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359-380. Academic Press, New York.
- E. Sandewall (1998). Cognitive robotics logic and its metatheory: Features and fluents revisited. *Electronic Transactions on Artificial Intelligence*, 2(3-4):307-329.