

Modeling Events, Actions, and Time

James F. Allen
Department of Computer Science
University of Rochester
Rochester, NY 14627

This brief note concerns what types of knowledge one must possess in order to be able to reason about events and actions. In particular, in comprehending stories or dialogues, many inferences are made based on what events and actions are described. These range from inferences about the temporal ordering of events to inferences concerning the beliefs and motivations of the actors. Here I will concentrate on the nature of events and actions and discuss their relation to temporal reasoning. The references below provide more detail on all these issues.

The formalism for actions and events used in most natural language understanding systems is based on case grammar. Each action is represented by a set of assertions about the semantic roles the noun phrases play with respect to the verb. Such a formalism is a start, but does not explain how to represent what an action actually signifies. If one is told that a certain action occurred, what can one conclude about how the world changed (or didn't change!). One possibility for such a mechanism is found in the work on problem-solving systems (e.g., [Fikes and Nilsson, 1971]), which suggests one common formulation of action. An action is a function from one world state to a succeeding world state and is described by a set of prerequisites and effects, or by decomposition into more primitive actions. While this model is extremely useful for modeling physical actions by a single actor, it does not cover a large class of actions describable in English. For instance, many actions seemingly describe non-activity (e.g., standing still), or acting in some non-specified manner to preserve a state (e.g., preventing your television set from being stolen).

Difficult problems also arise in this model concerning the simultaneous occurrence of actions in domains with more than one agent. For example, consider a simple blocks world with one block and two robots. Let there be two actions, PUSH_R, push the block to the right, and PUSH_L, push the block to the left. We would like to define the effect of these actions in terms of the block moving. But if the two robots perform a PUSH_L and PUSH_R simultaneously, the block does not move. Yet, we still want to say that each robot pushed the block. If we cannot express simultaneity of actions, the best we could do to model this situation would be to have the block oscillate as the robots pushed alternately.

The approach suggested here does not attempt to answer what an event or action actually is. Whatever an event is, the only way we can reason about one is by considering how the world changes (or remains constant) during some time interval in which the event occurred. Thus it is crucial that the temporal model in the logic be general enough to capture the scope of possible events. Actions are then defined as a subclass of events that involve agents and are described in a similar manner. The notions of prerequisite, result, and methods of performing actions do not play a central role in this study. While they are important for reasoning about how to attain goals, they don't play an explicit role in defining when an action can be said to have occurred. To make this point clear, consider the simple action of turning on a light.

There are few physical activities that are a necessary part of performing the action of turning on a light. Depending on the context, vastly different patterns of behavior can be classified as the same action. For example, turning on a light

usually involves flipping a light switch, but in some circumstances it may involve tightening the light bulb (in the basement), or hitting the wall (in an old house). Although we have knowledge about how the action can be performed, this does *not* define what the action is. The key defining characteristic of turning on the light seems to be that the agent is performing some activity which will cause the light, which is off when the action starts, to become on when the action ends. The importance of this observation is that we could recognize an observed pattern of activity as "turning on the light" even if we had never seen or thought about that pattern previously.

With this model, it is theoretically simple to describe two actions occurring simultaneously. The temporal conditions for each will be asserted to hold over the same time interval. It is then up to the reasoning component to infer any interactions that may arise. While this has not solved anything by itself, at least the complex problem can be expressed in the temporal logic, and reasoning techniques can then be investigated.

With respect to modeling time, I want to make just two basic claims. The first is that representations based on assigning dates for each time are unworkable. The second is that the underlying logic of time should be based on the notion of time intervals rather than time points.

There are many difficulties that arise in systems based on date lines. In such an approach, each time is represented by a value (e.g., a number) and relationships between times can be computed by some operation on the values (e.g., numeric ordering). One problem is that dates are not often supplied. Much temporal information in English is supplied only on a relative basis (e.g., E occurred before E'), both by the explicit mention of such relationships and by tense. For example, in the sentence

"We found the letter while John was away,"

the temporal connective "while" indicates that the time of the find event occurred during the time that John was away, and the past tense indicates that both events occurred in the past (i.e., before now).

The other major difficulty with date-based systems is that there can be considerable uncertainty in our temporal knowledge. For instance, we might know that either event E occurred before event E', or vice versa. But in any case, the times of E and E' did not overlap. One can only capture such information with a partial ordering relationship: no dates can be assigned that capture these constraints. This is not to say that dating is not a useful technique when it is possible, it just cannot be the foundation of the representation.

Turning to the time interval/time point controversy, we can easily observe that both appear to be referred to in English. Thus, we can say,

"We found the letter at 12 o'clock."
"We found the letter yesterday."

The most straightforward approach to dealing with time then seems to be to introduce points in time and then define intervals from those points (e.g., [McDermott, 1981; Bruce, 1972]). I do not use this scheme for two reasons. The first is

that such a representation is too uniform and does not facilitate structuring knowledge in a way convenient for typical temporal reasoning tasks. The second is that it encourages one to think of time as being isomorphic to the real line, which is a serious mistake.

The central issue concerning the first point is the importance of the *during* relation for reasoning. A major part of our temporal knowledge appears to be of the form

"event E' occurred during event E."

Our knowledge of the *during* relation allows a highly structured representation of time. In particular, a common way of inferring that some condition P holds during an interval T is to show that P holds in an interval that contains T. For instance, I might know that my office is locked today because it has been locked all week.

Furthermore, such a *during* hierarchy allows reasoning processes to be localized so that irrelevant facts are never considered. For instance, if one is concerned with what is true "today," one need consider only those intervals that are *during* "today," or above "today" in the *during* hierarchy. If a fact is indexed by an interval wholly contained by an interval representing "yesterday," then it cannot affect what is true now.

On the second issue, some annoying characteristics arise from allowing zero width of time points. For instance, two intervals that meet must either have a point in common or have a point between them. Thus to describe an event consisting of a light being transformed from being off to being on, either the interval where it is off meets the interval where it is on, and thus there is a point where the light is both on and off, or the interval where it is off is strictly before the interval where it is on, and thus there is a point between the two intervals where the light is neither on or off. This can be avoided by a technical trick such as treating all intervals as open on their beginning and closed on their end, but such tricks simply emphasize the unnaturalness of the approach. In an interval-based system, such issues need not arise: two intervals may meet without having any point in common.

Given this interval-based representation of time, what is the equivalent of time points? For instance, we often talk of the beginning or ending times of events. There is no reason to assume, however, that the beginning and ending times are instantaneous points. One might suggest that there is a minimum size ϵ of intervals, such that all intervals of size less than or equal to ϵ are considered to be points. The consequence of this would be that two such point intervals could then only be related by the relations $<$ and $=$. This approach is useful, but only if there is not one fixed value for ϵ , for the size at which an interval is considered to be a point depends on the reasoning task being done. For instance, the smallest time intervals we care about in everyday life are probably of the order of seconds, as physicists or computer scientists, we may consider times on the order of nanoseconds. Thus the interval size that we want to consider as points varies depending on the task as well as the proximity to the current time.

References

- Allen, J.F., "An interval-based representation of temporal knowledge," *Proc.*, 7th IJCAI, Vancouver, B.C., 1981.
- Allen, J.F., "What's necessary to hide?: Reasoning about action verbs," *Proc.*, 19th Annual Meeting, Assoc. Computational Linguistics, 77-81, Stanford U., 1981.
- Bruce, B.C., "A model for temporal references and its application in a question answering program," *Artificial Intelligence* 3, 1972.
- Fikes, R.E. and N.J. Nilsson, "STRIPS: A new approach to the application of theorem proving to problem solving," *Artificial Intelligence* 2, 189-205, 1971.
- McDermott, D., "A temporal logic for reasoning about processes and plans," Research Report 196, Dept. Computer Science, Yale U., March 1981.