Non-Standard Semantics for The Method of Temporal Arguments

Brian A. Haugh

Martin Marietta Laboratories
1450 South Rolling Road
Baltimore, Maryland 21227

Abstract

An expressive first-order temporal logic using the method of temporal arguments is developed and provided with a non-standard semantics that explicates its underlying temporal structure. Objections from the AI literature against the adequacy of temporal argument theories are answered in the course of discussing representational issues for the developed logic. This logic "accords a special status to time," distinguishes the temporal features of event occurrences from ordinary facts, and supports changing ontologies. We conclude that the method of temporal arguments remains a viable candidate for temporal reasoning in AI.

I Introduction

The "method of temporal arguments" (MTA) is our name for the approach to representing temporal information in which every ordinary predicate of a logic is supplemented by a special argument (or arguments) for time (e.g., *On(A, B, Time 1))*. This approach has been used for many years in database applications as the dominant means of representing the temporal aspects of database relations (e.g., Ann, 1986). Here, we show how to develop versions of this approach that have a more explicit temporal semantics and are capable of meeting the greater demands made on temporal representation and reasoning by AI systems. Previous criticisms of the suitability of the MTA approach are addressed in the course of discussing representational issues for such logics.

Temporal representation in AI (e.g., McDermott, 1982 and Allen, 1984) has been dominated by "reified temporal logics," in which terms referring to facts or propositions appear as arguments to truth predicates (e.g., *True(On(A, B), Time l))*. An interesting recent development in formalizations of these reified temporal logics has been the introduction of non-standard formal semantics for them (Shoham, 1986, and Reichgelt, 1986). These new formalizations provide special semantic foundations for explicitly identified temporal components of such logics, imparting greater clarity to their interpretation as well as a basis for meaningful proofs of soundness and completeness. Accompanying these developments have been several different arguments against the general suitability of the temporal arguments approach for AI systems. Here, we show how a non-standard semantics similar to those developed for reified temporal logics can provide clear foundations for a variety of different versions of the MTA approach, and how the objections against this approach may be answered, leaving it a viable candidate for temporal reasoning in AI. We begin by looking at reified temporal logics, and discuss the significance of their recent semantic developments.

II Reified Temporal Logics

Reified temporal logics have recently been defined (Shoham, 1986) as temporal logics in which propositions are made into objects (reified) by providing special propositional terms to refer to them. Temporal aspects of propositions are then indicated by asserting proposition types to be true (or to hold) at particular times. Thus, truth or holding predicates are used, taking propositional terms and temporal terms as arguments, to assert that the propositions hold during the specified times (e.g., *True(On(A,B), Tl))*. The best known examples of reified temporal logics, developed by James Allen (1983, 1984) and by Drew McDermott (1982), were formulated in ordinary typed first-order logics with no special formal semantics provided for their temporal features. Truth predicates were treated like any other predicate, and their temporal arguments, while distinguished by typing, were not given any special semantic treatment. As a consequence, no significant soundness or completeness results could be provided for the temporal aspects of the logic. Although the use of standard first-order logic provides formal assurance of soundness and completeness, it does not ensure that absurd temporal consequences cannot be derived, or that all of the genuine temporal consequences will be derived. A set of temporal postulates that allowed deduction of "B before A" from "A before B" would not affect formal soundness. Neither would an inability to deduce the transitive closure of simple temporal ordering relations affect the formal completeness of such a logic. Without a model theory that includes a distinctiy temporal component, any soundness or completeness results are inconsequential with respect to temporal reasoning.

While there have been clear temporal model theories for years for modal tense logics (e.g., Prior, 1967), difficulties in theorem-proving have hampered their implementation. Theorem-proving techniques with the elegance and completeness of resolution-based methods in first-order logic have not yet appeared for any quantified modal logics. Thus, the appearance of special temporal semantics for first-order theories offers the promise of the best of both worlds: a semantics as clear and explicit as that of the tense logics, with proof procedures as simple and elegant as those in ordinary first-order logic.

It should be noted, however, that, as yet, this promise remains unfulfilled, since soundness and completeness results have not yet been presented for these new logics. The absence of such results can be attributed largely to the novelty of this approach, but, may also be due, in part, to difficulties inherent in the reified temporal logics in which all such work has hitherto been pursued. There are special challenges in fully formalizing the propositional terms of reified temporal logics and the truth (or holding) predicates that take them as arguments, such as the well-known risks of inconsistency when using truth predicates in a first-order theory (Tarski, 1936). Similar risks are encountered in other uses of reified propositions, such as first-order theories of belief and knowledge (Montague, 1963 and Thomason, 1980). Only recently has it been shown that inconsistency may be avoided in such "self-referential" theories by careful restrictions on the propositional terms admitted by truth or belief predicates (Perlis, 1981, 1985).

Other difficulties arise for reified temporal logics in providing a satisfactory semantics for quantification into propositional terms, and in determining their scope (e.g., whether to allow embedding of truth predicates, temporal references, variables, logical connectives, or quantifiers in these terms). Early systems of reified temporal logics, (e.g., Allen's and McDermott's) provided no special formal syntax and semantics for their propositional terms. McDermott provided some axioms explicating his interpretation of propositions, but his formal semantics was standard FOL. Development of explicit syntactic restrictions on propositional terms and of special model-theoretic constructions for their semantics is one of the contributions of Shoham's pioneering work in providing better foundations for first-order temporal theories. Shoham's early reports of this work (Shoham, 1986) minimize the difficulties of handling propositional terms by restricting them to atomic predications. Reichgelt's subsequent work (Reichgelt, 1986) admits quantifiers and logical connectives to his terms, which provide greater expressive power, but at an unavoidable cost in complexity of syntax, semantics, and axiomatization. A significant body of work on representing propositional terms has also been pursued in association with investigations of logics for truth, knowledge, and belief (e.g., Perlis, 1981, 1985; des Rivieres and Levesque, 1986).

Eventually, the most capable knowledge representation systems will benefit from full capabilities for complex propositional terms, and quantification over and into them. Cause-effect relations with complex factual effects will require complex propositional terms to refer to them. Representation of belief and knowledge will also benefit from an ability to use belief predicates with arbitrarily complex propostional arguments. However, less ambitious systems may do well to avoid such complexity, and in such cases, the method of temporal arguments can provide a simpler alternative to reified temporal logics. An MTA approach is especially valuable for integration of knowledge bases with large temporal databases, since most (if not all) database applications of temporal representation involve some version of MTA. Furthermore, even when the fullest capabilities of propositional terms are required, the truth predicates of a reified temporal logic are unnecessary, since an MTA approach may be supplemented with propositional terms without recourse to truth predicates.

HI The Method of Temporal Arguments

The essence of the method of temporal arguments has already been defined as the use of a temporal argument (or arguments) in every predicate to establish temporal references. To distinguish it from the reified approach, we should note that it admits, as predicates, ordinary properties and relations (e.g., isred, isheavierthan) that are excluded by the reified approach. Even so qualified, this essence remains so minimal that it leaves room for enormous variety in particular versions. The basic logic is restrained only by the need for multi-place predicates, leaving many choices, such as: full, standard first-order logic (FOL), all sorts of restrictions on FOL, higher order logics, or any of a variety of non-standard logics. Temporal predicates for expression of temporal properties and relations, such as duration and ordering, can come in many types and combinations in a MTA logic. The basic temporal entities referred to by temporal arguments may be either points or intervals, and may be ordered in many different ways (dense, discrete, beginning, non-beginning, branching, ...). The most prominent alternatives for these choices are listed in Table 1. The more expressive alternatives are listed in boldface, when a clear advantage is present.

Much of the temporal representation in databases can be categorized under this approach (see, e.g., Ann, 1986), although database environments impose strong restrictions on logical capabilities (e.g., no disjunctions, no rules). The MTA approach can be very expressive when the more expressive features are chosen and supported by a full set of temporal predicates (e.g., *date(tl, dtl),* duration(t1,d1), *start(tl, t2), before(tl, t2),* ...).

1 Rule-based kb model	Relational db model
Variables & quantifiers	No variables
Full first-order logic	Horn clause logic
Dating	No dating
Ordering relations	No ordering
Branching order	No branching
Dense order	Discrete order
Beginning	No beginning
Ending	No ending
Durations	No durations
Interval times	No intervals
1 Point times	No points

Table 1. Alternative MTA Logics.

Here, we use a full first-order version, with a minimal set of ordering predicates, to illustrate the ease with which an explicit temporal semantics can be provided for a temporal arguments approach

IV Temporal Representation Issues

A variety of criticisms have been raised on the suitability of the temporal arguments approach. One broad criticism has been that it does not give any special status to time (Shoham, 1986). While this is true of standard FOL versions, the semantics provided below demonstrates that this special status can be easily achieved using techniques strongly analogous to those used by the new reified temporal logics. Other criticisms have involved the expressive advantages of reified propositions in causal relations, and problems in temporally dependent quantification. These criticisms are addressed below in the course of examining more general expressive issues that strongly influence the choice of crucial features of such logics. Discussion of these issues explains the motivation for many of the representational features chosen for our example MTA logic, in addition to answering the related criticisms.

A. Representing Times of Events

It is widely recognized that representing the times of occurrences of events requires a different sort of semantics than the times of the holding of ordinary facts. According to common usage, when a specific event occurs over a time interval T, it cannot properly be said to occur over any of the subintervals of T, or indeed, over any other interval. There is good reason for this convention as well. It supports the association of a unique time interval with each particular event, facilitating the identification and distinction of events as well as their countability. We may ask of event types, such as the throwing of a ball, how many times it has occurred, but it is not so reasonable to ask over how many time intervals an ordinary type of fact, such as a ball being red, holds. Ordinary durative facts typically hold over an uncountable number of time intervals since they hold over every subinterval of any interval over which they hold.

The generally recognized temporal nature of durative events requires that any logic capable of representing their times of occurrence be capable of representing temporal intervals, as well as distinguishing between facts and events. Thus, no traditional point-based tense logics (e.g., Rescher and Urquhart, 1971) can be used to reason properly about the times of event occurrences.

And, any fully expressive version of MTA will require a semantics capable of referring to intervals and of distinguishing between occurrences of events and the holding of ordinary facts.

Intervals may be represented in an MTA logic by a single temporal argument (e.g., *P(t, xl, ...xn))* taking interval constants and variables, by dual arguments representing the endpoints of an interval (e.g., *P(tl, t2, xl, ..xn)),* or by a single complex term, a function from two points into an interval (e.g., *P(f(tl, t2), xl, ..., xn)).* The last alternative was chosen here because it allows the use of different temporal functions to distinguish between intervals that are open on one end, open on both, or closed on both. The importance of this capability is discussed in the next section.

The time of occurrence of events in MTA can be expressed by an occurrence predicate (e.g., occurs(f(tl, *t2), event1)),* or simply by the times during which the case roles of an event hold (e.g., *agent(f(tl, t2), eventl, john)).* This latter approach is a natural temporal extension of a case frame representation of events, based on case grammars (Fillmore, 1968). The MTA approach is particularly well-suited for a temporal extension to case frames since the case frame roles (e.g., agent, activity, object, instrument, ...) can appear directly as top-level predicates relating the events to their case-role values at the time of event occurrence. Such predications about the case roles of events are clearly distinguished from the use of truth predicates and complex propositional functions (e.g., *True(tl, t2, stack(john, A, B))* found in reified propositional treatments, such as Shoham's (1986).

Case frame representations of events also have much to recommend them in addition to their ease of extension to temporal reasoning in MTA: modularization of event descriptions (Davidson, 1967); usefulness in natural language understanding and generation; elimination of rules required to determine case roles from functional representations; support of event tokens as well as event types; and efficiency of information retrieval via indexing of predicates and constant arguments in database systems and logic programming environments. A frame-based prototype of an MTA approach (Haugh, et.al., 1987) has been implemented using case frame representations of events within the TIMLS multi-representational knowledge engineering environment (Lewis, 1986), and used effectively for natural language generation of situation reports in a planning domain (Sekine, 1986). The logic presented here, however, will support either case frames or functions for representing events. A class of predicates is supported with the appropriate temporal features for event occurrences, and may be used in a variety of ways.

Some explanation on the advantages of terms referring to event tokens may be helpful here. Using names of specific events, rather than the event types typically used in reified temporal logics, allows asserting mat two events of the same type occur at the same time or overlapping times, and provides more direct support to the individuation and counting of events of a particular type. Exclusive use of event types would also create difficulties in determining that two different descriptions are descriptions of the same event, since individual events could not be named.

Despite these expressive advantages of using event tokens, exclusive use of event types would be no disadvantage in domains in which all the event descriptions were uniquely referring. In such circumstances, event types could be used exclusively in an MTA approach as easily as within reified temporal logics. Such an approach could be implemented in MTA within a case-frame formalism or using functional representations of event types. A functional representation of events within MTA would require the use of occurrence assertions (e.g., *occurs(f(tl,t2), throws(john, ball-1)),* whose forms might be indistinguishable from those of reified temporal logics. However, such a choice of representations would still be distinguished from a reified temporal logic, provided that events were clearly distinguished from propositions. Thus, we have shown several ways that an MTA approach may represent the time of occurrence and features of events without recourse to a reified temporal logic. A case-grammar-based representation supporting individual event names was chosen for its representational power, as well as its many general advantages.

B. Representing Points and Intervals

Besides the importance of representing time intervals, argued above, representation of time points is also very useful, especially for representing continuous variation of parameters over time. Although alternatives may be developed using discrete time or average values, the mathematics of continuous variation over points is much more straightforward and better understood. It is also convenient to be able to tie genuinely instantaneous events (e.g., electron orbital transitions), to a time point Time points also enable distinguishing between open and closed ends on intervals. If time points are to be represented for these reasons, and our earlier formulation of temporal references as intervals is to be used uniformly, then interval functions must be capable of representing points. This is easily done using an interval function that takes the same two points as arguments, returning the closed interval on them both (e.g., *closed(t1, t1)).* However, using closed intervals uniformly in our logic would create problems for representing certain kinds of relations between events. We could not have two events directly "meeting" one another (in Allen's terminology, Allen, 1984), if their times could be represented only by closed intervals. For events to meet, they must be defined over intervals that are open on one end and closed upon the other. This is the approach taken by Thomas Dean in his Time Map Manager (Dean, 1985), although genuine points cannot be represented thereby if it is used uniformly (hence, he makes use of infinitesimal intervals). To use interval functions to represent both time points and intervals supporting the meeting of events, two types of interval functions are needed, one closed on both ends and one half-closed. These capabilities are supported in the example logic by temporal functions (open, closed, open_l, open_r) that define all the endpoint variations on temporal intervals.

C. Distinguishing "Solid" and "Liquid" Facts

We have already argued the importance of distinguishing the temporal aspects of event occurrences from those of ordinary holdings of facts. This distinction closely parallels one from a set of such temporal distinctions developed by Shoham (1986). Predications of event occurrences are solid in Shoham's terminology because they are not true over any overlapping intervals. We use a bit stronger version of solidity here, in which solid predications are true only at a single unique time interval, since this is characteristic of event tokens. Shoham's liquid facts are the same as our predication class for ordinary facts wherein they hold over all subintervals. Our logic supports the typing of predicates according to their solidity, while Shoham reserves that distinction for axioms.

Although the capabilities we have described so far are adequate for events and ordinary facts, there are, arguably, predicates that cannot fit either category. Most action verbs seem to be of this type (e.g., solves(t, x, y), paints(t, x, y)), since they describe action types that may or may not be uniquely instantiated for any particular set of non-temporal arguments. Thus, a third category of predicates, used below for such relations, is not restricted to be wholly liquid or solid. Other categories could be defined analogously, if needed.

D. Representing Causation

Causation is one of those relations - like belief, knowledge, possibility and necessity - that can make good use of reified propositions when it is represented using predicates in a first-order logic. Expression of some event causing a fact to hold using causal predicates (e.g., *pcause(open_r(tl, tl), eventl, factl))* requires terms referring to facts. Thus, since reified temporal logics have already tackled the difficulties of referring to propositions, they are at some advantage. However, contrary to Shoham's suggestions (Shoham, 1986), it doesn't follow that a temporal arguments approach cannot achieve the same causal representation capabilities. Terms for propositions may be added to an MTA logic without adding any truth predicates, i.e., without converting to a reified temporal logic.

The simplest way to add propositional references to an MTA logic is to use a case-frame approach, analogous to that advocated for events above. Case grammars have a category for stative information that can represent simple atomic predications quite easily, with many of the same advantages noted above for case frame representations of events. With this approach, the logic may use ordinary individual constants and variables to refer to facts, and simple predicates for representing their case roles (e.g., verb, object, complement). This approach would provide much of the expressive power of ShohanVs logic, since his propositional terms are restricted to atomic predications.

More expressive fact causation, such as that developed by Reichgelt (1986), would require a more capable formalism, since case grammars were not designed to handle logical connectives and quantification. An embedded object language, such as Reichgelt's, appears the best approach to achieve the fullest generality, which will be required if belief or knowledge predicates are used in any case.

Thus, we can add abilities for representing causation and other relations involving simple reified propositions to our logic simply by providing predicates in the appropriate categories. A three-place causation predicate in the solid predicates category can represent causation, and a three-place belief predicate in the liquid category can represent belief. More complex propositions would require additional formal apparatus.

E. Representing Universal and Time-Specific Existence

An important fact about what exists (an ontology) is that it changes over time. This creates a problem for quantification in temporal logics, since we must be able to refer to everything that ever exists, as well as everything that exists at particular times. A shortcoming of many temporal theories has been the absence of support in the formal semantics for changing ontologies over time. Most versions of tense logics, for example, do not support this capability, nor do formulations of the method of temporal arguments. And, recently, it has been argued that the temporal arguments approach is inherently incapable of making this distinction (Reichgelt, 1986).

It would seem that the existence of an object at a certain time could be verified merely by determining if there are any predications that hold of it during that time. However, this test is inadequate, since many properties, like "is deceased," and relations, like "is the widow of," may hold at times at which the values of their arguments do not exist. Thus, Rescher and Urquhart, in discussing modal tense logics, have shown (Chapter XX, 1971) that they require either an additional quantifier or an existence predicate before they can express both kinds of existence. The method of temporal arguments can also achieve this expressive power when provided with such additional apparatus. In the MTA example below, a temporal existence predicate is used to formalize the existence of objects at particular times, and a single type of quantifier ranges over all objects existing at any time. Together, these features allow reference to just those objects that exist at all, or some, times during any identified time period.

V Special Semantics for MTA

The following example demonstrates an MTA-based logic whose special semantics satisfies the representational needs discussed in depth above. For simplicity of exposition, however, we have left out special predicates for dating, durations, and distances, since their needs are well-understood and their axiomatization is awkward.

A. Syntax

The primitive symbols consist of the elements of the following sets:

C a set of individual constants, {c1, c2, ...}

X a set of individual variables, {x1, x2, ...}

R1, R2, R3 sets of n-ary predicates, $\{P1^n_1, P1^n_2, \ldots\}$, $\{P2^n_1, P2^n_2, \ldots\}$, $\{P3^n_1, P3^n_2, \ldots\}$, for every integer $n > 0$

T a set of temporal constants, {T1, T2, ...}

N a singleton set containing the special temporal constant "NOW"

U a set of temporal variables, {t1, t2, ...}

TR a set of temporal predicates, {=, <, Exists}

TF a set of temporal functions, {open, closed, open_l, open_r}

L a set of logical symbols, { &, ¬, ∀,), (}

The set of temporal point terms, TP, consists of just the elements of T, N, and U. The set of temporal interval terms, TI, is defined as consisting of just the expressions of the form f(t1, t2), where f is a temporal function in TF, and t1, t2 are temporal point terms in TP. The well-formed formulae are then defined as follows:

1) If P is an n-ary relation in R1, R2, R3, t is a temporal interval term in TI, and $x_1, x_2, \ldots x_{n-1}$ are each individual terms from C or X, then $P(t, x_1, x_2, \ldots x_{n-1})$ is a wff.

2) If p and q are wffs, then (p & q) is a wff.

3) If p is a wff, then ¬p is a wff.

4) If p is a wff and x is a variable in X or U, then (∀x)p is a wff.

5) If t1 and t2 are each temporal point terms, then (t1 = t2) is a wff, and (t1 < t2) is a wff.

6) If tl is a temporal interval term, and x is an individual term in X or C, then Exists(t, x) is a wff.

B. Semantics

A model-structure <Times, Now, Order, Domain, Time_Domain> for the logic is defined as follows:

Times	is a non-empty set of time points.
Now	is a distinguished member of Times.
Order	is a binary ordering relation on the elements of Time.
Domain	is a non-empty set of individual objects, disjoint from Times.
TimeDomain	is a binary relation between subsets of Times and sets of elements of Domain

The ordering relationship for any specific logic would be furthur specified by meta-level axioms defining its characteristics (e.g., dense or discrete, branching or linear). A model is defined as a model-structure supplemented with an assignment function V that assigns:

a member of Domain to every constant in C

a member of Times to every constant in T

the element Now to the constant "NOW"

the equality relation on Times to "="

the Order relation to the predicate "<"

the TimeDomain relation to the predicate "Exists"

to the temporal function names "open," "closed," "open_1," and "openr," functions mapping from pairs of time points to sets of all the elements of Times between them, with "open" excluding both endpoints, "closed" including both endpoints, "open_1" excluding only the left (earlier), and "openr" excluding only the right endpoint

to every n-ary predicate, P, such that $P \in R1$, $P \in R2$, or $P \in R3$, a set of n-tuples of the form $<t, x_1, x_2, \ldots x_{n-1}>$, where t is a subset of Times, and each $x_i \in$ Domain.

The assignment to every predicate, P, in R2 is further restricted to provide the semantics of "liquid" facts by the requirement that if $<t1, x_1, x_2, \ldots x_{n-1}> \in V(P)$ and t2 is a subset of t1, then $<t2, x_1, x_2, \ldots x_{n-1}> \in V(P)$. The assignment to every predicate P in R3 is restricted to represent "solid" facts by the requirement that if $<t1, x_1, x_2, \ldots x_{n-1}> \in V(P)$ and t2 is not identical to t1, then $<t2, x_1, x_2, \ldots x_{n-1}>$ is not in V(P).

A variable assignment, G, is a function assigning each of the temporal variables in U some member of Times, and each of the individual variables in X some member of Domain. For ease of exposition, we consider G to extend the valuation V, when an interpretation is considered under a variable assignment. Truth under a variable assignment can be defined by further extending the assignment function, V, as follows:

1) For an atomic wff, where P is the predicate "Exists", or in R1, R2, or R3:
$V(P(t, x_1, x_2, \ldots x_{n-1})) = 1$ (true) iff
$<V(t), V(x_1), V(x_2), \ldots V(x_{n-1})> \in V(P)$;
otherwise $V(P(t, x_1, x_2, \ldots x_{n-1})) = 0$ (false).

2) For conjunction, if p and q are wffs:
$V((p \& q)) = 1$ iff $V(p) = 1$ and $V(q) = 1$;
otherwise $V((p \& q)) = 0$.

3) For negation, if p is a wff:
$V(\neg p) = 1$ iff $V(p) = 0$; otherwise $V(\neg p) = 0$.

4) For temporal ordering, if t1 and t2 are temporal point terms:
$V(t1 < t2) = 1$ iff $<V(t1), V(t2)> \in$ Order;
otherwise $V(t1 < t2) = 0$.

5) For temporal equality, if t1 and t2 are temporal point terms:
$V(t1 = t2) = 1$ iff $V(t1) = V(t2)$;
otherwise $V(t1 = t2) = 0$.

6) For quantification, if p is a wff, and x an individual variable in X or U:
$V((\forall x)p) = 1$ iff $V(p) = 1$ for all valuation assignments, G', like G except that x may be assigned any member of Domain or of Times, respectively;
otherwise $V((\forall x)p) = 0$.

Informally, atomic predicates of the form $P(t, X_1, x_2, \ldots x_{n-1})$ assert that some relation P holds between the objects $X_1, x_2, \ldots x_{n-1}$ during the time interval t. Predicates in R2 represent "liquid" relations, such as the colors of objects, those in R3 represent "solid" relations, such as event occurrences and causal relations, while those in Rl may be variable depending on their arguments. The existence predicate identifies the times at which objects exist. The distinguished constant "NOW" designates the present moment, allowing the past, present, and future to be distinguished.

VI Summary

We have shown how a clear semantics for the MTA approach can be created that "accords a special status to time." In the course of developing an MTA logic capable of representing the times of facts, events, continuous variation, and meeting of events, we have shown that at least two types of temporal intervals (closed and half-open) must be represented if interval functions based on points are used as temporal arguments. The use of event tokens instead of event types for representing the occurrence of events has been shown to be advantageous for individuating and counting events. A limited support of factual causation was shown to be supported by the basic formalism by use of a case-frame representation, without embedding the object language. Complete support of causation of arbitrary facts was shown possible without recourse to the truth predicates of reified temporal logic. Support of changing ontologies over time was incorporated into our example logic. Full axiomatization and soundness and completeness proofs have yet to be generated.

The central thesis supported by this exposition is the vitality of the method of temporal arguments as a clear foundation for temporal reasoning in AI. The explicit temporal semantics developed here provides the foundation for meaningful soundness and completeness results that were not available under the standard FOL interpretation. Such results, thus, establish a basis of confidence in temporal reasoning systems based on the MTA approach. In some contexts, the ease of implementation of efficient temporal reasoning systems using simpler versions of this approach, especially the ease of integration with temporal databases, provides advantages that no longer need be offset by an inadequate semantics. More complex representation systems, including embedded object language propositional terms, temporal persistence and clipping, and alternative possibility representations may be added to this firm foundation in a variety of ways. While the viability of the MTA approach has, thus, been vindicated, no decisive advantages have been found either for it or for reified temporal logics in their most expressive versions.

Acknowledgements

I would like to thank John Lewis, Stuart Pearlman, Donald Perlis, Han Reichgelt, Yukiko Sekine, and Frederick Suppe for helpful comments on earlier versions of this paper.

References

Ahn, I. 1986. " Towards An Implementation of Database Management Systems with Temporal Support." *In Proc: International Conference on Data Engineering,* IEEE Computer Society Press. pp.374-381.

Allen, J. 1983. "Maintaining Knowledge about Temporal Intervals." *Communications of the ACM,* 26:11: 832-843.

Allen, J. 1984. "Towards a General Theory of Action and Time." *Artificial Intelligence,* 24:2: 123-154.

Davidson, D. 1967. "Causal Relations/' *Journal of Philosophy,* 64:691-703.

Dean, T. 1985. "Temporal Imagery: An Approach to Reasoning about Time for Planning and Problem Solving." Ph.D. Thesis, Department of Computer Science, Yale University.

des Rivieres, J. and Levesque, H. 1986. "The Consistency of Syntactical Treatments of Knowledge.'' *In* Halpern, J. (Ed.) *Theoretical Aspects of Reasoning About Knowledge: Proc. 1986 Conference,* pp. 115-130.

Fillmore, C. J. 1968. "The Case for Case." *In* E. Bach and R. T. Harms, (eds.), *Universals in Linguistic Theory.* New York: Holt, Rinehart and Winston, pp. 1-88.

Haugh, B. and Lewis, J., et.al. 1987. "TIMLS: Logic Based Tools for Planner Knowledge Bases." Martin Marietta Labs Tech. Report (forthcoming).

Lewis, J. 1986. "The Inference Machine Laboratory: Graphic Tools For Knowledge Management." *In Proc. IEEE Graphics 86, Vancouver, B.C.*

McDcrmott, D. 1982. "A Temporal Logic for Reasoning about Processes and Plans." Cognitive Science, 6: 101-155.

Montague, R. 1963. "Syntactical Treatments of Modality, with Corollaries on Reflexion Principles and Finite Axiomatizability." *Acta Philosophica Fennica.,* 16:153-167. Also in Thomason, R.H. (ed.) 1979. *Formal Philosophy: Selected Papers of Richard Montague.* New Haven: Yale University Press, pp. 268-301.

Perlis, D. 1981. "Language, Computation, and Reality." Ph.D. Thesis, University of Rochester.

Perlis, D. 1985. "Languages with Self-Reference I: Foundations." *Artificial Intelligence* 25:301-322.

Perlis, D. 1986. "Self-Reference, Knowledge, Belief and Modality." *AAA1-86,* pp. 416-420.

Prior, A.N. 1967. *Past, Present and Future.* London: Oxford University Press.

Reichgelt, H. 1986. "Semantics For Reified Temporal Logic." Department of A.I. Research Paper No. 299, University of Edinburgh.

Rescher, N. and A. Urquhart. 1971. *Temporal Logic.* New York: Springer-Verlag.

Sekine, Y. 1986. "TIMLS/MUMBLE Text Generation Project." Martin Marietta Labs Tech. Report 86-94.

Shoham, Y. 1986. "Reified Temporal Logics: Semantical and Ontological Considerations." *In Proc. 1986 European Conference on Artificial Intelligence.* Society for the Study of Artificial Intelligence and Simulation of Behavior (ALSB), pp. 390-397.

Tarski, A. 1936. "Der Wahrheitsbegriff in den formalisierten Sprachen." *Studia Philos.,* 1:261-405.

Thomason, R. 1980. "A Note on Syntactical Treatments of Modality." *Synthese* 44:391-395.