The Knowledge Engineering Review, Vol. 11:3, 1996, 281–288

Trends in temporal representation and reasoning

LUCA CHITTARO and ANGELO MONTANARI

Dipartimento di Matematica e Informatica Università di Udine Via delle Scienze, 206-33100 Udine, Italy (email: {chittaro/montana} @dimi.uniud.it)

1 Introduction

Time is one of the most relevant topics in AI. It plays a major role in several of AI research areas, ranging from logical foundations to applications of knowledge-based systems. Despite the ubiquity of time in AI, researchers tend to specialise and focus on time in particular contexts or applications, overlooking meaningful connections between different areas. In an attempt to promote cross-fertilisation and reduce isolation, the Temporal Representation and Reasoning (TIME) workshop series was started in 1994. The third edition of the workshop was held on May 19–20 1996 in Key West, FL, with S. D. Goodwin and H. J. Hamilton as General Chairs, and L. Chittaro and A. Montanari as Program Chairs. A particular emphasis was given to the foundational aspects of temporal representation and reasoning through an investigation of the relationships between different approaches to temporal issues in AI, computer science and logic.

This paper is not so much a report of the workshop, as an analysis of the major trends in temporal representation and reasoning in AI as they emerged from the presented papers. It is organised into two main sections. The first section witnesses recent developments in traditional themes, such as temporal constraints and reasoning about actions and events. The second one pinpoints emerging trends, such as time granularity and temporal reasoning in medical informatics.

The space devoted to each research topic is roughly proportional to the resonance it had at the workshop. Before citing any new result, we provide a preliminary discussion to introduce the reader to the specific topic, and to put the new contributions in the proper perspective.

2 Established lines of research: recent developments

In this section, we focus on reasoning about temporal constraints and reasoning about events and actions, which are commonly recognised as the two main lines of research in the field of temporal reasoning. We give a short overview of the basic approaches and results obtained in these two areas, and outline some recent developments. Since temporal logics are the backbone of reasoning about events and actions, we add a subsection on recent developments about them and their execution. We conclude the section by briefly discussing the problem of dealing with uncertainty in temporal knowledge.

2.1 Temporal constraints

Mechanisms for reasoning about temporal constraints are concerned with the handling of temporal relations on the basis of the properties of the underlying temporal domain (ordering, boundedness, metric, etc.) and make no assumptions about the sorts of facts that can be said to hold over points or intervals of time. Temporal relations can thus be interpreted as constraints on time-dependent entities.

Reasoning about temporal constraints is performed using a number of different algorithms for constraint satisfaction. Nevertheless, it can be abstractly described in terms of two basic functionalities, namely, simplification and propagation. *Simplification* returns the minimal equivalent subset of a given set of constraints on a finite set of temporal variable (points or intervals). *Propagation* produces the *deductive closure* of a set of constraints, looking for all constraints that can be derived from the originally given ones. The combined application of propagation and simplification makes it possible to compute the *minimal representation* of the deductive closure of a set of constraints. This representation consists of all and only the relevant, globally consistent constraints, and can include both initially given constraints as well as constraints derived by propagation. Moreover, if the initial set includes incompatible constraints, the computation returns the empty constraint. In such a way, it can be used to verify the *consistency* of any set of constraints. The minimal representation of a set of constraints can also be used to generate some/all possible assignments of values to variables and to verify if a given assignment is consistent with the given constraints. A formal characterisation of the simplification and propagation functionalities, together with some examples of their use for solving temporal reasoning problems, is given in Montanari and Pernici (1993).

Several formalisms for expressing and reasoning about temporal constraints have been proposed in the literature. They can be classified either according to the type of temporal constraints they deal with (qualitative, quantitative, or both) or according to the type of temporal elements they consider (intervals, points, or both). Some relevant formalisms for dealing with qualitative temporal constraints are Allen's *Interval Logic* (1983) (IA), Vilain and Kautz's *Point Algebra* (1986) (PA), which is as expressive as the proper subset of IA (called SIA) of *pointisable interval relations*, and Nebel and Bürckert's *ORD-Horn subalgebra* (1995). The set of relations in the ORD-Horn subalgebra strictly includes the relations of SIA, but it does not include the disjointness relations of IA. Quantitative temporal constraints are dealt with by Dechter, Meiri and Pearl's *Distance Algebra* (1991) (DA), which models distances between points as well as durations of intervals. Finally, Meiri (1991) defines a formalism combining quantitative and qualitative temporal constraints.

It goes without saying that there is a trade-off between expressiveness and computational complexity of the various algebras. The problem of computing the minimal representation of a set of constraints or reporting its unsatisfiability is brought back to the problem of computing the minimal consistent labelling of the corresponding network. Computing such a labelling as well as finding one or all possible models of the network for IA requires an exponential time in the worst case. The same results hold for DA when no restrictions to constraint types are imposed. Approximations of the exact solutions for IA and DA networks can be obtained in polynomial time using consistency algorithms. They remove local inconsistencies that could never be part of a global solution, but they do not always compute the minimal label between any two nodes, and thus they do not guarantee global consistency. One, two and three-consistency (generally referred to as node, arc and path consistency, respectively) are the most commonly used algorithms.

ORD-Horn subalgebra is the largest tractable subalgebra of IA. Exact solutions can thus be obtained in polynomial time for the ORD-Horn subalgebra and for all less expressive algebras. The complexity bounds of these algorithms, however, are still unacceptable for domains in which a large database of relations needs to be managed (*scalability problem*). In recent years, *ad hoc* algorithms have been developed to take into account the specific temporal structure underlying relevant classes of problems (e.g. in domains such planning and story understanding temporal data tend to fall naturally into chain-like structures) in order to significantly improve performance.

Schwalb and Dechter (1996) extend the work on quantitative temporal constraint satisfaction problems (TCSPs) by proposing different ways of weakening path consistency which improve significantly on the algorithms proposed in Dechter, Meiri and Pearl (1991). A problem with enforcing path consistency with quantitative TCSPs is the so-called fragmentation problem, i.e. some intervals may be broken into several smaller subintervals with an increase in the number of intervals per constraint that may result in an exponential blow up. To solve this problem, Schwalb and Dechter (1996) propose (different variants of) an algorithm, called *Loose Path Consistency*, based on a simple intersection operator between sets of intervals.

Wetprasit, Sattar and Khatib (1996) show how to represent recurring events as multi-point events in a suitable extension of PA. They also propose an algorithm for finding feasible relations for multi-point event networks, and identify the special cases for which it can find the exact solution.

Lamma, Mello and Milano (1996) face the problem of efficiently implementing temporal constraint solving in a logic programming framework. They exploit the constraint handling mechanisms of *finite domain constraint logic programming* (CLP(FD)), proposing a two-module architecture that deals with both qualitative and quantitative temporal constraints.

2.2 Temporal logics

Several logic-based formalisms have been proposed and extensively used for representing temporal knowledge in AI. Traditionally, they were roughly partitioned on the basis of the paradigm (interval-based vs. point-based) they choose. More recently, several research efforts have been devoted to the integration of the paradigms as well as to the integration of temporal and spatial knowledge. An increasing attention has also been paid to the application of temporal logic to classical tasks such as planning and monitoring. Temporal logic and its application in planning are the subjects of van Benthen (1996) and Bacchus (1996), respectively.

The complexity of temporal logics is the main focus of van Benthem (1996). On the basis of the observation that the complexity of temporal phenomena is often a reflection of their formal modelling as much as of their intrinsic nature, he discusses several logical strategies for remodelling temporal logics so as to lower their complexity. Then, he introduces a two-sorted geometrical temporal model, having points/states and (finite and infinite) paths as primitives, and shows how this model makes it possible to overcome traditional oppositions, such as point-based vs. interval-based temporal ontologies. This allows one to view the various temporal logics proposed in the literature as specific theories of states and paths. Finally, he presents two possible ways of hierarchically ordering temporal logics. On the one hand, there exists a natural syntactic hierarchy of modal fragments, starting from Prior's tense logic, running upward to 'Since'/'Until' logics, and beyond. On the other hand, a matching semantic hierarchy of two-sorted point-path bisimulations can be introduced that relates temporal models that are geometrically alike to some degree.

Bacchus (1996) shows the usefulness of temporal logic for planning. While traditional planning concentrated on primitive actions and instantaneous states, the introduction of a temporal dimension allows one to reason about sequences of actions and/or about the sequences of states they generate. The temporal logics he adopts are quite simple: their syntax contains the usual temporal modalities and their models are linear sequences of states. Nevertheless, they can be successfully used to specify search control information as well as to specify goals that are temporally extended over time, that is, to constrain the way the agent achieves its final goal. Furthermore, he shows how temporal logic can be used to define history-dependent rewards to be assigned to states in order to generate optimal plans.

Traditional approaches to the problem of modelling common sense physics have often dealt with temporal and spatial aspects in isolation. Gagné and Trudel (1996) describe a spatio-temporal ontology suitable for representing and reasoning about the qualitative motion of rigid bodies, with the goal of supporting a uniform treatment of motion in one, two, and three dimensional spaces, and provide such an ontology with a succinct axiomatisation in first-order logic.

There are essentially two possible approaches to the problem of executing temporal logics: the direct approach of finding a proof algorithm for a specific class of temporal logics; and the translation of temporal logics into classical first-order logic to use predicate calculus mechanical theorem prover to implement temporal theorem provers. The problem of directly and efficiently executing temporal logic is addressed in Dixon (1996). The paper introduces a breadth-first search resolution-based procedure for Propositional Temporal Logic, and states its soundness and completeness. A comparison with the previously developed depth-first search algorithm is also given. Becher (1996) presents a class of first-order modal temporal logics that extend the interval temporal logics proposed in Halpern and Shoham (1991) with the notion of generalised intervals.

Instead of developing complex direct methods for executing the considered logics, the author defines a translation function that maps them into a standard two-sorted first-order logic, preserving satisfiability and validity.

2.3 Reasoning about actions and events

Reasoning about the consequences of a given set of actions and/or events is fundamental in many reasoning tasks such as planning and plan validation. This kind of temporal reasoning, called *temporal projection*, essentially fills in the gaps of incomplete temporal data by drawing conclusions on the basis of incomplete information. Information which is neither explicitly asserted nor monotonically implied by the available knowledge can be inferred, according to suitable assumptions, such as closed world, default persistence, default change, and so on. Conclusions derived in this way are *defeasible*, that is, they can be withdrawn if the addition of new information makes them inconsistent. This requires the addition of nonmonotonic features to temporal systems, e.g. the replacement of classical negation with negation as failure.

Derivation of temporal information from incomplete temporal data is performed by means of forward and backward temporal projections with respect to an arbitrary time point ("reasoning as of"). *Forward temporal projection* concerns the prediction of the future, and can be seen as the process of inferring consequences of specific facts on the basis of what is currently believed, using general domain knowledge. *Backward temporal projection* concerns the explanation of the past (post-diction), and can be seen as the process of inferring justifications of specific facts on the basis of what is currently believed, using general domain knowledge. Justifications may consist of new links among already known facts as well as of new facts. More precisely, *forward temporal projection* supports inferences of the form: "if X is true at this time then Y is true at that *no earlier* time", while *backward temporal projection* supports inferences of the form: "if X is true at this time then Y is true at that *no later* time". Forward and backward projections are based on cause-and-effect relationships defined among temporal entities. Cause-and-effect relationships provide a model of change that relates the validity of effects to the validity of their relevant causes. The degree of confidence of inferences drawn using forward and backward temporal projections is different. Forward projection leads from causes to their *certain* effects, while backward projection leads from effects to their *possible* causes.

The Situation Calculus (McCarthy and Hayes, 1969) and the Event Calculus (Kowalski and Sergot, 1986) are the most well-known formalisms for reasoning about actions and events. The ontology of the Situation Calculus is mainly based on the concepts of situation and action. An action is the cause of transition between two situations, which are periods of time during which there are no changes in the properties of the represented domain. The two main differences between the Event Calculus and the Situation Calculus are that (i) the Event Calculus introduces a real line of time points, and properties may or may not hold over each point, and (ii) it has no explicit notion of situation: events happen at time points, and initiate and/or terminate intervals over which *specific* properties hold. Sandewall (1994) proposes a systematic approach to the representation of knowledge about dynamical systems. The *features and fluents* framework he developed aims at formally characterising the range of applicability of the different nomonotonic logics of actions proposed in the literature, that is, it allows one to associate with any such logic the class of reasoning problems for which the logic obtains exactly the intended conclusions. Sandewall's framework, however, only deals with actions that occur sequentially.

Yi (1996) provides an extension of Sandewall's features and fluents framework to take into account concurrent actions with duration. However, he restricts himself to the case of actions that do not change each other's course.

Mohktari and Kayser (1996) introduce a causal theory based on a Kripke-like possible worlds semantics, focusing on how the effects of voluntary actions may vary according to the context (e.g., they define an inhibit function that returns the events which can inhibit the result of an action).

Vila and Schwalb (1996) consider the problem of representing both discrete (e.g., a digital signal) and continuous (e.g., a linearly decreasing quantity) change in a uniform framework. To this end, they provide a general theory of time that diversifies fluents according to whether the change in the parameter they model is continuous or discrete.

2.4 Uncertainty in temporal knowledge

Representing and reasoning about uncertainty in temporal knowledge aims at dealing with events and actions whose effects cannot be univocally determined (e.g., developing an illness does not necessarily imply that all the effects and complications of that illness will occur). To this end, logic-based formalisms which assign a boolean truth value to events and/or properties at each time point or over each time interval are inadequate.

To suit the needs of uncertainty in temporal representation and reasoning, the adopted formalism should allow one: (i) to replace boolean truth values by probabilities, which express the certainty degrees of events and/or properties; (ii) to model the indeterminacy in the relationships between an event and its possible effects: an event produces a given effect with a given probability; (iii) to take into account the influence of the context on the effects of an event: the probability that an event produces a given effect can vary according to contextual conditions, and the probability of the different conditions must be considered. Moreover, temporal aspects themselves often play a role in determining probabilities. For example, the effects of a fever depend also on how long it lasted, and the probability of a fault in an engineered device is higher when it is new (due to manufacturing defects) or when it is old (due to wear, stress, aging, etc.).

Researchers working at integrating time and probability typically extend temporal approaches by introducing mathematical machinery such as Markov processes and Bayesian networks. In this context, Tawfik and Neufeld (1996) concentrate on the issue of irrelevance. The problem is to identify which information is relevant or not in order to draw a conclusion, and relevance is influenced by time, e.g., if a cat was seen in the living room two weeks ago and we are interested in its actual location, the previous information is not of much help (on the contrary, if the cat has been seen two minutes ago, the relevance greatly increases). The proposed approach models the degeneration of relevance with time by using a Markov chain representation.

3 Emerging trends

Besides the previously discussed evolution of established areas of research, temporal representation and reasoning is both expanding itself towards other fields (e.g., databases), and growing new subareas (e.g., time granularity). This section discusses some of the most significant developments in these directions.

3.1 Time granularity and temporal abstraction

A relevant case of incomplete temporal information, which is not dealt with by standard temporal reasoning systems, is provided by time information specified at different time granularities. *Time granularity* can be defined as the resolution power of the temporal qualification of a fact. Providing a formalism with the concept of time granularity makes it possible to specify time information with respect to differently-grained temporal domains. This does not merely mean that one can use different time units—say, months and days—to represent time quantities in a *unique flat temporal model*, but it involves more difficult semantic issues related to the problem of assigning a proper meaning to the association of facts with the different temporal domains of a *layered temporal model* and of switching from one domain to a coarser/finer one. Mechanisms to handle time granularity consist of default projection rules to map a fact from the granularity level with which it is originally associated into another granularity level. Projection from coarser to finer granularities and from finer to coarser ones are performed by means of downward temporal projection and upward temporal projection, respectively. *Upward temporal projection* supports inferences of the form: "if X is true at this time of a given domain then Y is true at that time of a *coarser* domain", while *downward temporal projection* supports inferences of the form: "if X is true at this time of a given domain then Y is true at that time of a *finer* domain".

Abstraction synthesises single higher-level concepts from rough time-stamped data. In this way, it allows one to abstract from the considered domain only those aspects which are relevant to the actual goal. As an example, in a clinical context where a huge number of patient data is regularly acquired, an emerging pattern over a span of time has much more significance than an isolated finding or even a set of findings. According to Shahar (1996), the temporal abstraction task can be decomposed into five parallel subtasks: (i) *temporal context restriction* creates the relevant contexts for interpretation of data, crucial for focusing and limiting the scope of the inference; (ii) *vertical temporal inference* computes values of higher-level concepts starting from values of contemporaneous input data or abstractions; (iii) *horizontal temporal inference* computes values of concepts at a given level starting from similar-type propositions that hold over different time intervals at that level; (iv) *temporal interpolation* bridges the gaps between similar-type, but temporally disjoint, propositions to create longer intervals; (v) *temporal pattern matching* creates intervals by matching patterns over disjoint intervals over which propositions of various types hold.

It is worth noting that time granularity and abstraction are *orthogonal* dimensions, that is, changing the level of time granularity does not necessarily require to change the level of abstraction, and *vice versa*.

Providing a formalism with the notions of time granularity and abstraction increases both its expressive and computational power. With respect to the *expressive power*, time granularity and abstraction increase both the temporal distinctions that a language can make and the distinctions that it can leave unspecified. This means that considering two events as simultaneous or temporally distinct, or two time dependent relations as temporally overlapped or disjoint, depends on the time granularity one refers to. Similarly, we can analyse the evolution of a quantity at the detailed level of sampled measurements as well as at the abstract level of qualitative states, e.g., only distinguishing stable and unstable states, or normal and abnormal states. Temporal abstraction is affected by the context, e.g., the concept of normal/abnormal state of a patient may vary according to the specific pathology. With respect to the *computational power*, time granularity and abstraction support different grains of reasoning to deal with incomplete and uncertain knowledge. They also allow one to tailor the visibility of the knowledge base and the reasoning process to the needs of the actual task. Secondly, they allow one to alternate among different levels of time granularity and/or abstraction during the execution of a task in order to solve each incoming problem at a level of time granularity and/or abstraction as coarse as possible. Finally, they also possibly allow one to solve a problem at a level of time granularity and/or abstraction coarser than the required one. Such a simplification speeds up the reasoning, but implies a relaxation of the precision of the solution.

The ability of providing and relating temporal representations at different "grain levels" of the same reality is widely recognised not only as an important research theme, but also as a major requirement for many AI applications, such as medical knowledge based systems.

Mota and Robertson (1996) propose a framework for representing and reasoning about the behaviour of interacting agents/processes which work at different time granularities, and show how it can be used to represent and reason about *ecological systems*. In particular, they show how the underlying temporal structure, consisting of a hierarchy of time-cycles, allows one to easily represent natural events which are inherently cyclic and hierarchically related. Cukierman and Delgrande (1996) propose a general approach to the problem of representing and reasoning about schedulable, repeated activities, specified using calendars. Repeated activities must be taken into consideration in scheduling systems, financial systems, process control systems, and, in general, *date-based systems*. The paper points out a set of parameters according to which different types of temporal repetitions can be classified.

Finally, Combi, Pinciroli and Pozzi (1996) emphasise the relevance of time granularity in the *clinical domain*, and propose a data model (TIME-NESIS) aimed at appropriately storing and managing anamnestic information given by the physician at various and mixed levels of time granularity.

3.2 Temporal databases

There is an increasing evidence that a number of important problems in AI have a natural counterpart in relational databases, and *vice versa*. As an example, there exists a close relationship between constraint satisfaction problems in AI and the satisfaction of join-dependencies in relational databases. Thanks to this relationship, decomposition techniques developed for relational databases can be successfully applied to decompose constraint satisfaction problems. The existence of similar connections between problems in the area of temporal representation and reasoning, and in the area of temporal databases has been investigated in recent years.

Unlike conventional databases, which are designed to store the most recent available data, temporal databases maintain both the current and the past states of the database. Moreover, each state is not a simple snapshot of the reality, but a temporally extended description of it, involving past, present and possibly future data. Since in temporal databases many of the data change their values over time, it makes sense to provide them with the machinery for performing basic temporal inferences. In particular, many temporal databases give the possibility of managing temporal information not only to retrieve information as stored in the database, but also to automatically derive further data. To this end, temporal databases must take into account that temporal information may be incomplete (events may be only partially ordered), metric and layered (events may be associated with different time granularities and their occurrence times may be known within metric bounds) and defeasible (assumptions about the persistence of certain facts may turn out to be unwarranted as new information is acquired).

Bettini, Wang and Jajodia (1996) present a general framework for describing time granularity (database) systems. They identify the main parameters according to which the various systems can be differentiated, and formally characterise the relations among granularities in these systems. Then, they introduce the notion of network of temporal constraints with several granularities, point out the semantical and computational differences from constraint networks with a single granularity, and show that the associated consistency problem is NP-hard. Finally, they present a polynomial approximate algorithm.

Orgun (1996) considers a temporal relational algebra based on a temporal relational data model in which the database is modelled as a set of time-varying relations. Each time-varying relation is viewed as a collection of ordinary relations indexed by time points. The collection of time points is modelled as the set of natural numbers. The resulting algebra is a pointwise extension of the relational algebra over the set of time points. As temporal logic, it keeps time implicit, and hence it does not allow its explicit manipulation. Orgun (1996) proposes an extension of such an algebra with two new constructs: tagging and recursion. The resulting algebra supports recursive queries through the addition of an equational component.

3.3 Time in problem solving

The handling of time in problem solving is often faced with the problem that general temporal reasoning systems proposed in the literature are difficult to tailor to specific applications. In most cases, this leads to the development of *ad hoc* solutions that cannot be reused even in very similar applications. To overcome this situation, recent research focuses on the design of general techniques for dealing with temporal aspects of specific tasks such as qualitative simulation and medical monitoring.

Brajnik and Clancy (1996) consider the task of qualitative simulation, i.e., constructing a set of possible behaviors consistent with a model of a dynamical system represented by qualitative versions of differential equations. They extend the well-known qualitative simulator QSIM introducing a propositional linear-time temporal logic, to provide temporal constraints on the trajectories of variables as an additional input to the simulator. Then they enforce the validity of these temporal constraints during simulation in order to restrict the set of possible evolutions as opposed to testing the outcomes of the simulation after its completion.

A particular mention has to be given to the medical informatics area, where much work is currently being devoted to the design of temporal problem solving methods. For example, the already mentioned temporal abstraction mechanisms proposed by Shahar (1996) are being used for abstraction of clinical data in the patient monitoring task. Keravnou (1996) describes the requirements for temporal reasoning in medical knowledge-based systems and discusses their realisation in terms of a generic temporal kernel to be embedded in such systems, emphasising the relevance of reasoning at multiple granularities. Bichindaritz and Conlon (1996) focus on the temporal component of case-based reasoning in medicine, where time is needed to appropriately describe the history of patient cases used for concept learning purposes.

References

All the 1996 references appear in: Chittaro L, Hamilton H, Goodwin S, Montanari A. (eds.), *TIME-96: Proceedings of the Third International Workshop on Temporal Representation and Reasoning*, IEEE Press, 1996.

Allen, JF, 1983. "Maintaining knowledge about temporal intervals". *Comm. ACM* **26** (11) 832–843.

Dechter, R, Meiri, I and Pearl, J, 1991. "Temporal constraint networks". *Artificial Intelligence* **49** 61–95.

Halpern, JY and Shoham, Y, 1991. "A propositional modal logic of time intervals". *J. ACM* **38** (4) 935–962.

Kowalski, R and Sergot, M, 1986. "A logic-based calculus of events". *New Generation Computing* **4** 67–95.

McCarthy, J and Hayes, PJ, 1969. "Some philosophical problems from the standpoint of artificial intelligence". In: Meltzer, B and Michie, D. (eds.), *Machine Intelligence 4*, Edinburgh University Press.

Meiri, I, 1991. "Combining qualitative and quantitative constraints in temporal reasoning". *Proc. of AAAI-91*, Anaheim, CA, 260–267.

Montanari, A and Pernici, B, 1993. "Temporal reasoning". In: Tansel *et al.* (eds.), *Temporal Databases: Theory, Design and Implementation*, Benjamin/Cummings.

Nebel, B and Bürckert, HJ, 1995. "Reasoning about temporal relations: a maximal tractable subclass of Allen's Interval Algebra". *J. ACM* **42** (1) 43–66.

Sandewall, E, 1994. *Reasoning about Change: Time and causation from the standpoint of Artificial Intelligence*, Oxford University Press.

Vilain, MB and Kautz, H, 1986. "Constraint propagation algorithms for temporal reasoning". *Proc. of AAAI-86*, Philadelphia, PA, 377–382.