Aggregation and Forgetting: Two Key Mechanisms for Across-Time Reasoning in Patient Monitoring

M. Dojat
INSERM U.296, Faculté de médecine,
8, av. Gral Sarrail 94010, Créteil France
dojat@laforia.ibp.fr

C. Sayettat
UTC, Centre de Recherche de Royallieu
BP 649 60206, Compiègne Cedex, France
sayettat@hds.univ-compiegne.fr

Abstract

In this paper, we describe the methodology we used to model temporal aspects of a reasoning for interpreting clinical data in real-time.

We assume that 1) aggregation of similar observed situations and 2) forgetting non relevant information are two essential mechanisms involved in this class of reasoning processes. These two mechanisms require domain and context specific knowledge. There were gradually introduced by refinement of initial domain-independent temporal entities and inferences.

We implemented our temporal model in order to integrate it into a working prototype, mainly reactive in its reasoning, for the management of ventilator therapy, and we are in the process of testing it in an Intensive Care Unit (ICU).

1. Introduction

The need for an explicit representation of time in medical diagnosis has been advocated by several researchers. More generally, all systems that are supposed to interact with realistic worlds are time-varying. Thus, many formal studies in artificial intelligence relating to time and continuous change representation have been performed [McDermott 1982; Allen 1983; Kowalski and Sergot 1986; Galton 1990; Van Beek 1992]. Unfortunately, there is a big gap between this advanced work and the integration of a time map manager into real-time systems (such as intelligent monitoring systems) essentially due to the complexity of the temporal constraint propagation algorithms, and due to the expressive power insufficiency of the techniques proposed compared to the richness of the situations to be modelled.

For real applications, domain-dependent considerations have to be introduced into temporal reasoning systems to ensure computational tractability [Williamson and Hanks 1993]. We adopted a different kind of approach than that based on constraint satisfaction methods.

Our work is an endeavour to mimic the physician's attitude who tries to interpret dynamically physiological data. We propose a model using two fundamental mechanisms: *aggregation* and *forgetting*, appropriate for a class of applications such as real-time patient monitoring.

In this paper we present our model: temporal ontology, structures and relations, and temporal inferences. Then, we show how we implemented this model in order to integrate it into a working prototype for the management of ventilator therapy and finally we will discuss some aspects of this work relating to other approaches.

2. Temporal ontology

Our ontology divides the world between "existing" domain entities that by essence have no temporal dimension (description of the world), and temporal entities that are time stamped and are used to develop a temporal discourse about the changing world. We consider the recognition of change, which is context-dependent, as the central point in the perception of time.

The two primitive entities introduced are *Event* and *State*. An event is an instantaneous object associated with a time instant that often causes change in our world representation. In contrast, a state characterises a local property of the world that lasts over a time period. As proposed initially by [Kowalski and Sergot 1986], the occurrence of an event *initiates* or *terminates* a state and a state is assumed to persist until an event terminates it. Thus, it is valid over a period defined by events corresponding either to the start or to the end of a time period. We consider that accurate time-stamps for the clinical observations are available, thus temporal structures are temporally totally-ordered.

3. Temporal inferences

In process control, driven by the collected data the expert must 1) diagnose the current state of the system, 2) build an interpretation of the system's behaviour, 3) predict the system's evolution with regard to previous states and if necessary, 4) construct and execute a plan of actions to drive the system to an expected state. Across time, many successive states could be diagnosed and accumulated.

Guided by a natural necessity to act, the expert must *aggregate* similar states, recognise only relevant changes, and *forget* non-relevant information, in order to build a global dynamic interpretation of the system's behaviour. Some temporal abstractions are processed to modify the length and the position of a mobile temporal window that brings to light a minimal set of states, representative of the actual situation and useful for the current reasoning. All states not visible through the window are forgotten, at least for a while. In the patient clinical course, *aggregation* and *forgetting* take place at several abstraction levels and their activation is context-dependent.

4. Implementing temporal model

The model proposed has been applied to represent time and change occurring in a medical expertise for management of patients who suffer from a lung disease and are mechanically ventilated.

4.1. Medical context

The clinician has to adapt the level of the mechanical assistance to the evolutive physiological needs of the patient. His/her reasoning is based on the interpretation of time periods passed in given situations. He/she must quickly react in the case of a sudden interruption in a continuous state of correct ventilation. Depending on the context (such as patient with a low level of assistance or patient with stable ventilation for a long period), he/she must tolerate some instabilities without modifying the current action plan.

To construct a global interpretation of the patient's ventilation, he/she must forget tolerable instabilities and group similar ventilatory states [Dojat and Sayettat 1993]. This concise representation is used to adapt his/her action plan: a gradual decreasing of the assistance to progressively re-educate respiratory muscles in the case of a global correct ventilation; a rapid increasing of assistance for a persistent incorrect ventilation.

4.2. Temporal entities

We simulated our medical universe using an object-oriented approach and chose the Smalltalk-80 system to represent all the entities (non temporal and temporal objects and behaviours) that we needed. We modelled the way concepts belong to one another using taxonomic links among entities. A temporal metric was available with time points (dates) and intervals. All operations on events (such as comparison) and states (such as Allen's relations) were defined with message-passing. From initial domain-independent temporal entities hierarchy, we introduced new domain-dependent temporal objects to model the temporal part of the expertise. In our application, diagnosis, actions on ventilator and alarms were treated as events; current ventilation (normal, insufficient ...), predicted situations (stable, aggravated...) and expected situations (restoration, weanable...) as states.

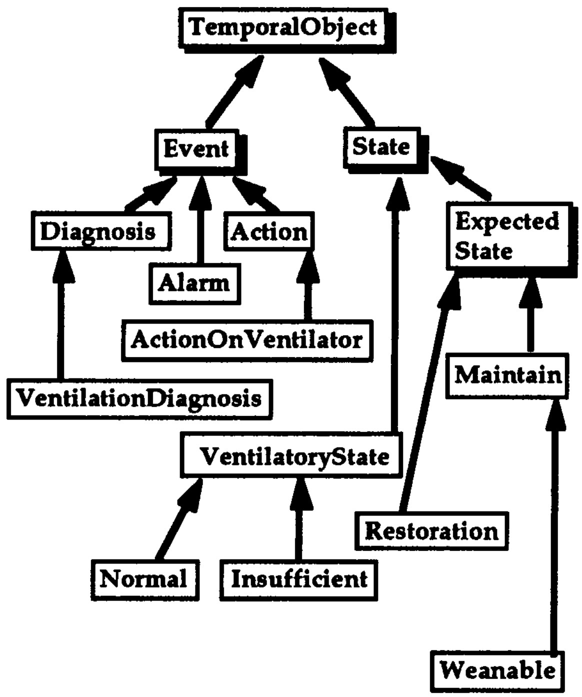

Figure 1: Excerpt from the hierarchical organisation of temporal entities. Shadowed rectangles represent domain-independent temporal objects. Arrows indicate a Kind-Of relation.

4.3. Temporal inferences

Event-state relationships such as *initiates* or *terminates* were represented as methods for instances of `Event` class so they can be understood by all temporal objects that may have an effect on our world representation. All relations reflecting notions of persistence or break were transformed in methods for `State` instances.

Aggregation and Forgetting are the two keys mechanisms for across-time reasoning we have studied. They were modelled with first order production rules thanks to NéOpus [Pachet 1992], a rule-based programming system embedded in Smalltalk-80. The natural combination of forward-chaining rules and object-oriented programming so obtained, give us a powerful tool to implement our model. Class and rule-base inheritance allowed us to gradually refine the initial domain-independent aggregation and forgetting rules to introduce domain specific semantic considerations.

Let see the general rule R1, for partial continuity, which was implemented in the rule-base *TemporalAbstractions* at the top level of the temporal rule-bases hierarchy. This domain-independent rule considered a state (s1) as a continuous state if the duration of the meeting observed instability (s2) does not exceed a value and is followed by a state (s3) similar to s1.

```
partialContinuity

| State s1 s2 s3. Duration d1 |
s3 persistent.
s1 sameAs: s3.
s2 contradicts: s3.
s2  between: s1 and: s3.
s2 duration <= d1.
actions
s1 increaseValidityOf: s3 duration.
s2 forget. s3 forget.
```

The general rule R1

This rule was redefined in a rule-base subclass of *TemporalAbstractions* to take into account of domain-dependent knowledge. When a patient has been judged as a good candidate for the weaning (a special strategy to gradually decrease the assistance in order to withdraw the machine), several interruptions due to incorrect ventilation are tolerated (rule R2) before restoring a low level of assistance, considering anew the patient as weanable and pursuing the initial action plan. R1 matches all State instances and R2 matches only instances of `Weanable` class.

```
partialContinuity
| Weanable w1 w2 |
w1 broken.
w1 discontinuityIsTolerable.
w1 < w2.
actions
p2 forgets.
p1 increaseValidityOf: p2 duration.
```

The specific rule R2

Note that the recognition of change was represented by two methods of `State` class `sameAs:` and `contradict:` . They were redefined in subclasses to indicate the degree of change required to consider two instances as similar or in contradiction.

Rules are compiled into a Rete network that contains at each inference cycle all the candidates to be matched by a rule. Inferences are applied to all objects present in the network. On the contrary, the expert utilises some knowledge to reason with a limited number of information. Thus, we represented two types of forgetting: active and passive forgetting to remove non relevant objects from the Rete network (equivalent to short term memory) and send them into long term memory. *Active forgetting* consisted in statements in the action part of some rules (see R1 and R2 actions). *Passive forgetting* (repression) was represented with meta-rules that allow a declarative control of the deductive process. Passive forgetting meta-rules were fireable under general conditions: for example when the number of objects overpasses a limit (overload) or under specific conditions: for example when a state is persistent for a long time, previous situations are not more considered by the expert. As aggregation rules, forgetting rules were redefined at several levels in the inheritance tree for rule-bases and meta-bases.

5. Discussion

The approach we propose aims at introducing an explicit time and change representation into a real-time patient monitoring system while ensuring computational tractability. To build our model, we elicited, from practical medical expertise in ICU, two mechanisms used by experts: aggregation and forgetting. Several researchers [Allen 1983; Kohane 1987; Ash and Hayes-Roth 1990] have proposed mechanisms similar to aggregation to avoid computational burden of temporal reasoning.

Aggregation and forgetting depend on semantic considerations and triggering conditions have to be adapted dynamically to the reasoning strategy. In our approach, they are high level mechanisms inherent to the clinician's reasoning that operate on symbols extracted from clinical data.

The aggregator Peptide proposed by [Weld 1990] detects automatically repeating cycles without knowing in advance entities equivalencies as it is the case in our approach. However, it was designed as an exploring tool for a future theory of aggregation and its applicability to aggregate information under real-time conditions is not reported.

Our reflection is closer to the recently published work of [Shahar and Musen 1993] with the RéSUMé system that performs temporal abstractions of time-stamped data for patient monitoring. RéSUMé is designed for long-term therapies management where disease's progress is monitored episodically. Thus, the authors have introduced facilities to revise their temporal conclusions when new information about the past is available or to interpolate data and appreciate gradual change over long periods. For a system that receives data continuously and whose the goal is to administrate a therapy in real-time, a model of patient's evolution as a state transition problem and a representation of change as a monotonic discrete process seem more adapted and efficient.

6. Conclusion

The proposed temporal abstractions allow clinical observations to be incrementally interpreted as they are acquired. We have integrated the temporal package into NéoGanesh [Dojat and Pachet 1992], a closed-loop control system mainly reactive in its reasoning, used for ventilator management, and we are in the process of testing it in ICU. We should like to extend and re-used our system in similar medical applications. To reach this goal, new data sets typical of a critical care environment will be useful to complete the validation.

Acknowledgements

This work was supported in part by Hamilton Medical AG.

References

McDermott, D., *A temporal logic for reasoning about processes and plans.* Cognitive science, 1982. **6**: 101-155.

Allen, J.F., *Maintaining knowledge about temporal intervals.* Comm. ACM, 1983. **26**, (11): 832-843.

Kowalski, R.A. and Sergot, M. J., *A logic-based calculus of events.* New generation computing, 1986. **4**: 67-95.

Galton, A., *A critical examination of Allen's theory of action and time.* Artificial Intelligence, 1990. **42**: 159-188.

Van Beek, P. *Reasoning about qualitative temporal information.* Artificial Intelligence, 1992, **58**: 297-328.

Williamson, M. and Hanks, S. *Exploiting domain structure to achieve efficient temporal reasoning..* In proceedings of the 13th IJCAI, 1993, 152-157.

Dojat, M. and Sayettat, C. *Temporal reasoning in a medical expertise.* In proceedings of the 5th SEKE, conference, 1993. San Francisco, 134-141.

Pachet, F., *Représentation par objets et par règles : le système NéOpus.* 1992, PhD thesis, Paris VI.

Kohane, I., *Temporal reasoning in medical expert systems.* 1987, PhD thesis, MIT.

Ash, D. and Hayes-Roth, B. *Temporal representation in blackboard architectures.* Technical report KSL 90-16, Stanford University, Calif, 1990.

Weld, D. S. *The use of aggregation in causal simulation.* Artificial Intelligence, 1990. **30**: 1-34.

Shahar, Y. and Musen, M. A. *RéSUMé: a temporal-abstraction system for patient monitoring.* Computers and biomedical research, 1993. **26**: 255-273.

Dojat, M. and Pachet F.. *An extendable knowledge-based system for the control of mechanical ventilation.* In proceedings *14th IEEE-EMBS* conference, 1992. Paris, 3: 920-921.