


LABORATORY FOR
COMPUTER SCIENCE



MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

MIT/LCS/TR-389

**TEMPORAL REASONING
IN MEDICAL EXPERT SYSTEMS**

Isaac S. Kohane

May 1987

This blank page was inserted to preserve pagination.

Temporal Reasoning in Medical Expert Systems

by

Isaac Samuel Kohane

April 1987

© Massachusetts Institute of Technology 1987

This research was supported by National Institutes of Health Grant No. R01 LM04493 from the National Library of Medicine and National Institutes of Health Grant No. R24 RR01320 from the Division of Research Resources.

Temporal Reasoning in Medical Expert Systems

by

Isaac Samuel Kohane

Abstract

Diseases develop and change over time. Much of the distinction between pathophysiological complexes rests on the temporal relations of their component events. Therefore, knowledge bases that fail to capture the temporal component of the course of disease omit useful diagnostic knowledge. Expert systems that cannot reason with temporal knowledge are impaired in distinguishing between hypotheses and therefore have to explore much larger problem-spaces than would a human or temporally sophisticated expert system. Temporally naive expert systems are also limited in the extent to which they follow human diagnostic style and provide reasonable automated explanations and diagnostic questions.

The Temporal Utility Package (TUP) is a domain independent utility that is designed for use with a wide variety of knowledge representations. TUP can represent points, intervals, qualitative and quantitative temporal relations, groups of points, common temporal "yardsticks," and alternate temporal contexts. TUP employs a form of constraint propagation to make temporal inferences. As the inference computation grows rapidly with the number of points, TUP enables temporal deductions to be performed locally by "chunking" the temporal data base. The knowledge structures of the application domain can be used to automatically guide this "chunking" process. Certain aspects of TUP's performance may have their parallel in human cognition.

THRIPHT is a prototype expert system that demonstrates TUP's application and the role of temporal reasoning in different phases of the diagnostic process: data gathering, hypothesis evocation, elaboration, instantiation, and hypothesis ranking.

TUP and THRIPHT together illustrate why temporal reasoning is necessary for successful second generation medical expert systems, and how to provide this capability.

Keywords: temporal reasoning, medical diagnosis, expert systems.

Acknowledgments

- My parents, Akiva and Judith, for the love, moral support, and personal attention over twenty-seven years without which none of this work would have been possible.
- My brother, Danny, who knows how and when to kickstart me.
- Peter Szlovits and Bill Henneman for their patience, guidance, friendship, and support, even when I was unaware of my own ignorance.
- The graduate students at the Clinical Decision Making Group at M.I.T. for the camaraderie and lunchtime discussions. In particular, Mike Wellman whose criticisms led me to many insights, Elisha Sacks for patient explanations, Tom Russ for introducing me first to ITS and then Lisp Machines, and Phyllis Koton for being an exemplary and entertaining officemate.
- Barry Shein, system wizard *par excellences*, and Mary Riendeau for providing a congenial and reliable environment for the development of TUP and THIRPHT.
- Ramesh Patil and Bill Long for solid advice.
- Robert Friedman, Michael Klein, and Ruth Levine, at the Boston University Medical Center, who were always available and yet were tolerant of my own unavailability.
- Dan Bernstein, for his help as friend, career counselor, and personal physician.

This report is a modified version of my PhD thesis [28].

Contents

Abstract	i
Acknowledgments	ii
List of Figures	vi
List of Tables	vii
1 Introduction	1
1.1 The Patient History	2
1.1.1 Time and the Patient History	3
1.1.2 Organization of the Diagnostic Task	6
1.2 Technical Goals	11
1.2.1 Breadth of Representation	11
1.2.2 Computational Feasibility	12
1.2.3 Interaction of Temporal and Atemporal Reasoning: Domain Independence	12
1.3 Results	13
1.3.1 Expert System Demonstration	14
1.4 Organization	15
2 The Temporal Utility Package	17
2.1 The Range Relation	17
2.1.1 Diagrams and Descriptions	18
2.1.2 Examples	19
2.2 Generating Temporal Inferences	27
2.2.1 Range Addition	27
2.2.2 Constraint Propagation	29
2.2.3 Temporal Dependencies	31
2.2.4 Contradiction Handling	33
2.2.5 Tools for Studying Constraint Propagation	34
2.3 RREL Retrieval	37
2.4 Restricting Constraint Propagation	38
2.4.1 Implementation of Clustering: The Reference Set	40
2.4.2 Costs of Clustering	42

2.4.3	Computational Savings	43
2.4.4	Design of a Temporal Clustering Heuristic	43
2.5	Automatic Generation of Reference Sets	45
2.5.1	Using Other Knowledge Representations	48
2.5.2	Qualitative Reasoning	49
2.5.3	Planners	51
2.5.4	Frame-based and Hybrid Representation Languages	51
2.6	Reference Systems and Reference Sets	54
2.6.1	The PRESENT Reference Set	57
2.7	Predicates and Retrieval Functions	58
2.7.1	Predicates	58
2.7.2	Event Retrieval	59
2.8	Persistence	61
2.9	Unresolved Issues	61
2.9.1	Reasoning With Temporal Uncertainty	62
2.9.2	Recurring Events	64
2.9.3	Parameterized Bounds	66
3	Application of TUP to Medical Diagnosis	68
3.1	Data Collection	69
3.2	Hypothesis Evocation	71
3.3	Elaboration	79
3.3.1	Defining Mutually Exclusive Courses	80
3.3.2	Generating Contexts	81
3.3.3	Generating Reference Sets	82
3.4	Instantiation	82
3.4.1	Consequences of Constraining a Hypothesis Context	83
3.5	Ranking Patient Models	85
3.5.1	Evaluating Hypotheses	85
3.5.2	Differentiating Between Patient Models	86
3.6	Annotated Example	87
4	Temporal Knowledge Base Engineering	98
4.1	Implicit "Common Sense"	98
4.1.1	"Common Sense" Temporal Knowledge	98
4.2	Adherence to TUP Semantics	99
4.2.1	RREL Bounds	99
4.2.2	Loose Bounds and Alternate Contexts	100
4.3	Volume of Relevant Information	102
4.4	Example	102

CONTENTS

v

4.5	Finding Reference Sets	104
4.6	Growth in the Number of Contexts	105
4.7	Summary	105
5	Time In Human Cognition	107
5.1	Extent of Temporal Computation	107
5.2	Constraint from Background Temporal Information	110
5.3	Relative Performance of Retrieval Tasks	110
5.4	Summary	112
6	Conclusion	113
6.1	The Problem	113
6.2	The Plan	114
6.3	The Outcome	115
6.4	Complications	117
6.5	The Prognosis	118
A	Tables and Summaries	119
A.1	Predicates and Retrieval Functions	119
A.1.1	Predicates	119
A.1.2	Assertions	120
A.1.3	Point Functions	120
A.1.4	Relation Retrieval	120
A.2	Range Addition	122
A.3	Interval to Range Relation Conversion	123
A.4	Range Relation to Interval Relation Conversion	127
B	RRELS of the Annotated Example	130
B.1	RRELS of a Subhypothesis of Hepatitis B	130
B.2	Notes on the RRELS	132
B.3	Temporal Assertions from the Patient History	132
B.4	Notes on the Assertions of the Patient History	132
C	Abbreviations	134
	Bibliography	135

List of Figures

1.1	Misdiagnoses of Hepatitis B	7
1.2	Phases in the Taking of a Patient History.	8
2.1	Time-Line and Graph Form of Temporal Diagram	20
2.2	Illustration of assertion of example 2.	21
2.3	Illustration of assertion of example 3.	21
2.4	Illustration of assertion of example 4.	23
2.5	Illustration of assertion of example 6.	25
2.6	Illustration of assertion of example 7.	26
2.7	Temporal Data Base After Range Addition.	28
2.8	Initial Constraint.	30
2.9	First propagation of constraint.	30
2.10	Final propagation of constraint.	31
2.11	Propagation of the "OVERLAPS" Assertions.	32
2.12	Constraining the "OVERLAPS" assertion.	32
2.13	Loop Formation	36
2.14	Effect of Reference Set Configuration upon Constraint Propagation.	41
2.15	Triggered hypothesis for hepatitis B	47
2.16	Process Description of Fluid Flow.	50
2.17	Use of Reference System Information.	56
2.18	Patterns of P Waves and QRS Complexes in ALTERNATION.	65
3.1	Part of the Tangled Hierarchy of Triggers.	78
3.2	The Patient History	90
3.3	The elaborated hypothesis for hepatitis B	93
3.4	Triggered hypothesis for non-A, non-B hepatitis	94
3.5	Instantiated Hypothesis Without Patient History	96

List of Tables

A.1	Range Addition	122
A.2	Interval Relation To RREL Conversions	123
A.3	RREL to Interval Relation Conversions	127
C.1	Abbreviations	134

1. Introduction

In 1982, Drew McDermott [38] wrote,

A common disclaimer by an AI [Artificial Intelligence] author is that he has neglected temporal considerations to avoid complication. The implication is nearly made that adding a temporal dimension to the research (on engineering, medical diagnosis etc.) would be a familiar, but tedious exercise that would obscure the new material presented by the author. Actually, of course, no one has ever dealt with time correctly in an AI program, and there is reason to believe that doing it would change everything.

There are many reasons why this should be. Phenomena that we associate with time's passage, such as change, are ubiquitous. This is reflected in our view of the world, and the language and notation we use to describe it. It should therefore not be surprising if expert systems, bereft of a conceptual representation and vocabulary of time, as well as the necessary mechanisms to make temporal inferences, fail in general to interact smoothly with us. Expert system performance and knowledge acquisition is at best incomplete and at worst fatally flawed if the system does not possess a systematic and principled "understanding" of temporal knowledge. Therefore, even though temporal reasoning is not in and of itself sufficient to produce satisfactory behavior in expert systems, it is a necessary and significant step in that direction.

In contrast to the state of affairs in the late 1970's, described by McDermott, in the past six years several projects have addressed the problem of incorporating temporal reasoning in AI programs. The emphases of these projects vary with the application domain. Some areas of knowledge seem to be more dependent upon and demanding of temporal knowledge than others. The initial "blocks world" experiments [65] were adequately served by a simple ordering of events. Other applications require considerably more sophistication. Medical diagnosis and therapy are often selected as test tasks, not so much because they provide realistic tests of the breadth of representation and the generality of the temporal reasoning mechanism, but because without such capabilities, performance in these domains is poor.

Within the domain of medical expert systems, the efforts in temporal reasoning have tended to focus on issues most relevant to intensive care. That is, domains

in which there are large amounts of real-time data available, and where the values of several physiological variables are known and can change rapidly and often. In such applications the emphasis is on abstracting point data, correcting erroneous information, and making predictions about future values based on perceived trends [16,32,33]. In contrast, the temporal issues involving the task of patient history-driven diagnosis, have only been cursorily examined. This task is the focus of this thesis. In the sections that follow, I justify this choice and briefly outline the specific temporal issues involved.

1.1 The Patient History

Apart from serving as an investigational tool to study the issues of temporal reasoning, patient history-driven diagnosis is a worthwhile task in terms of the broader objectives of AIM [Artificial Intelligence in Medicine]. The patient history remains the most powerful and widely used diagnostic tool of health-care practitioners. It serves to narrow the diagnostic possibilities so that other, much more expensive, diagnostic modalities can be employed sparingly. Often, the patient history is the *only* available means for establishing the diagnosis.¹ Consequently, a significant amount of educational resources is spent to augment and refine the clinician's ability to obtain and use the patient history. Capturing and automating such widely applied expertise, will therefore provide health-care providers with an intellectual tool with which to manage many of the problems of the medical "information explosion."

Increasingly, clinical practices—especially large hospitals—are attempting to bring patient records, including the patient history, on-line [2,4,62]. The availability of such data holds the promise that expert systems, running in the "background", will be able to use the patient histories to perform diagnostic tasks without requiring special attention or effort from health personnel. To be able to effectively use such a patient history, it is necessary that the expert system be capable of sophisticated temporal reasoning as is illustrated below.

¹For instance, in distinguishing bronchiolitis from asthma, or differentiating rheumatoid arthritis from systemic lupus erythematosus, the pattern of disease the patient reports might be the only available clue.

1.1.1 Time and the Patient History

Any pathophysiological process is a creature of many dimensions, one of which is time. The patient history is a window onto this creature. If this window filters out temporal information, then errors will be made in describing the process and distinguishing it from others. Moreover, it is often inadequate to substitute ordering, whether partial or complete, for quantitative specification of temporal distance, as this quantification may be essential to characterize the process. In order to provide a flavor of the different forms of temporal expression in a patient history, a transcript of a fictitious patient visit is presented in the next section. Following this, I discuss the role of temporal reasoning in the taking of the patient history.

Scenario

The patient, a young male (J.), his mother (Mrs. W.) and the physician (Dr. C.) are in the physician's office. The mother reports concern with "fussiness, rapid-breathing, cough, and loss of appetite". After the preliminary discussion:

- Dr. C.: "How old is J.?"
- Mrs. W.: "Two years old."
- Dr. C.: "When did he first become 'fussy'?"
- Mrs. W.: "This past Sunday, when he woke up, around 8 a.m."
- Dr. C.: "Was he coughing that morning?"
- Mrs. W.: "No, that started later, seven to eight hours later. In the early afternoon."
- Dr. C.: "When did you first notice that J. was breathing rapidly?"
- Mrs. W.: "In the morning, two days later."
- Dr. C.: "When did J. stop eating, or stop eating as much as he usually does?"
- Mrs. W.: "Monday."
- Dr. C.: "A day before you noticed his rapid breathing?"

- Mrs. W.: "No, during the same day he started to breathe rapidly... Tuesday, I suppose."
- Dr. C.: "Anybody in the family have a cold?"
- Mrs. W.: "Everyone did, J.'s younger sister too — but she had bronchiolitis."
- Dr. C.: "Is the rest of the family well now?"
- Mrs. W.: "Yes, my husband was last to get it. He's felt fine for over a month."
- Dr. C.: "Has J. ever had anything like this before?"
- Mrs. W.: "Yes, now that you mention it, it happened about the same time last year. The doctor at the time said it could be bronchiolitis or asthma. But it cleared up after J. had spent one night in the hospital."
- Dr. C.: "Anybody in the family have asthma?"
- Mrs. W.: "Yes, both my husband and my brother-in-law had asthma when they were children, but neither have had any attacks since."

Upon further work-up, the physician strongly suspects that asthma is the correct diagnosis. The patient's response to therapy supports this suspicion. The patient and his mother are seen one month later for follow-up. During this visit, the following conversation arises.

- Mrs. W.: "Is J. going to keep having these attacks?"
- Dr. C.: "A child like J., diagnosed as having asthma, may follow several clinical courses. Some young people with asthma will continue to have attacks during their childhood. Of those, a minority have asthma in their adult years. Most probably, your son will be symptom-free by the end of his adolescence—as was the case for your husband and brother-in-law. Meanwhile, we can manage this problem without restricting J.'s activities."

There are many temporal concepts in each of the above queries and assertions. Among these:

- References to common temporal yardsticks such as the calendar, stages of development (e.g. in infancy, childhood and adolescence) or numerical age.
- Relative positions of events, both quantitative (e.g. coughing "seven to eight hours" after beginning to behave "fussily.") and qualitative (e.g. anorexia "during" the period of tachypnea).
- Temporal specification with respect to the present (e.g. "He's felt fine for over a month." which is an implicit specification of the period immediately prior to the present.)
- Alternate temporal hypotheses. In the above example, the duration of the period in which the asthma would be symptomatic depended on the clinical course the patient followed (i.e. whether the disease would resolve with the onset of adulthood).

Expert System Performance

The import of temporal representation is not restricted to the range of expression of an expert system and the "human-like" quality of the man-machine dialogue. It also makes the performance of the expert system considerably more efficient and focused. To illustrate this, take the three findings of jaundice, abdominal pain, and blood transfusion. Imagine that these are three (as in figure 1.1(a)) of very many findings in the patient history in an automated medical record. A temporally naive expert system would have to include the hypothesis of transfusion-borne acute hepatitis B high in the differential diagnosis. It might be however, that the temporal relationships in this case were such that to a human medical expert, there would never be any question, or at the most a low likelihood of this diagnosis. The jaundice may have happened during the neonatal period, the abdominal pain could have happened after an appendectomy, and the transfusion may have occurred during a later caesarian section.

Now, if we were to replace the previous, temporally naive expert system with one capable of causal reasoning, then the expert system would be able to use the temporal precedence relations between causal antecedent and consequent to distinguish between some hypotheses. If such an expert system were given a patient history that included the assertions that a blood transfusion preceded both abdominal pain and jaundice (as in figure 1.1(b)), then the hypothesis of transfusion-borne acute hepatitis B would again rank high in the differential diagnosis. However, if

a physician had been told that the transfusion had preceded the jaundice by only one day, a lot of other causes for the jaundice would come to mind, because the incubation period of the hepatitis B virus is considerable longer than one day.

Were we to go one step further, to enable the previous expert system to represent and reason with quantitative temporal information, this would still be inadequate. For instance, if such an expert system were given a patient history that included the assertions that the jaundice had followed the blood transfusion by two months, and the abdominal pain had followed the transfusion by 45 days (as in figure 1.1(c)), then yet again the program would rank acute transfusion-borne acute hepatitis high in the differential diagnosis. If however, the jaundice had occurred 20 years ago, then for the human expert, the diagnosis would not be a leading element of the current differential diagnosis.

The point made by these examples is that the ability to use a large variety of temporal information permits a dramatic pruning of the problem space—the number of hypotheses to be considered. This is important both because it lightens the computational burden of the expert system, and also because it cuts down requests for patient information and tests that might involve extra financial costs and unnecessary patient morbidity. This pruning of the problem space also has the secondary effect of producing a more focused and human-like diagnostic style, as many obviously unreasonable hypotheses are almost immediately excluded from serious consideration.

1.1.2 Organization of the Diagnostic Task

The examples above illustrate some of the different genres of temporal *expression* that can be found in a patient history. To discuss the various roles of temporal *reasoning* in patient history-driven diagnosis, I have arbitrarily, divided this task into five distinct phases. As described below, these phases may appear to follow each other in linear sequence, but as diagrammed in figure 1.2, there may be multiple, cyclic paths through these phases.

Initial Data Gathering

The first phase of the diagnostic task is that of data-collection, or history-taking. The activity in this phase is largely driven by the data collected, rather than by any hypothesis of the patient's pathophysiological status. This is to be distinguished from the later phases of the diagnostic task, after the evocation of hypotheses,

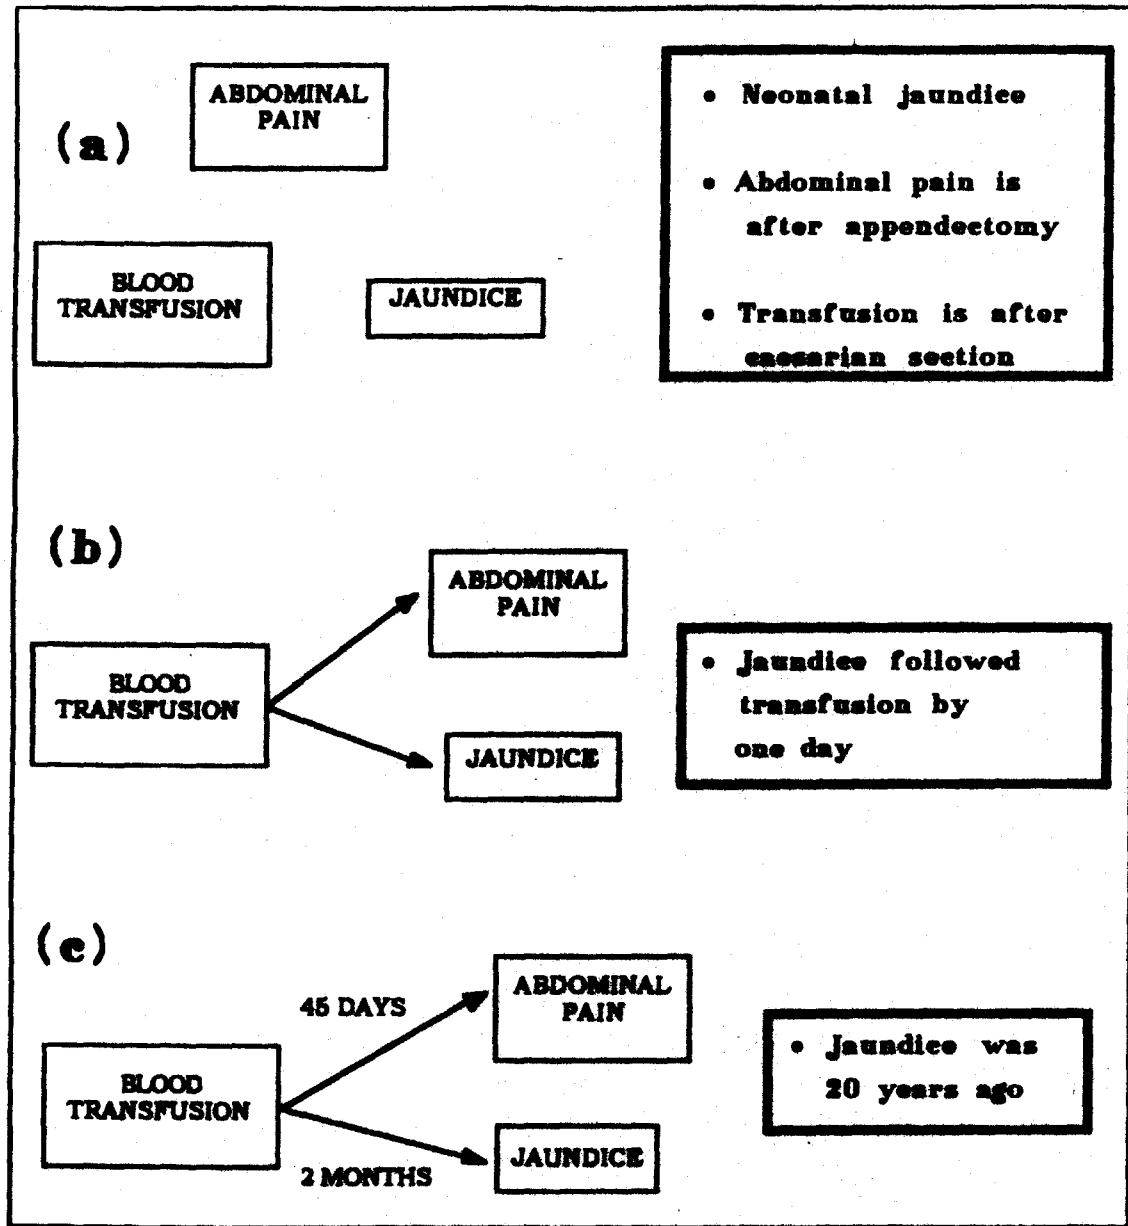


Figure 1.1: Misdiagnoses of Hepatitis B

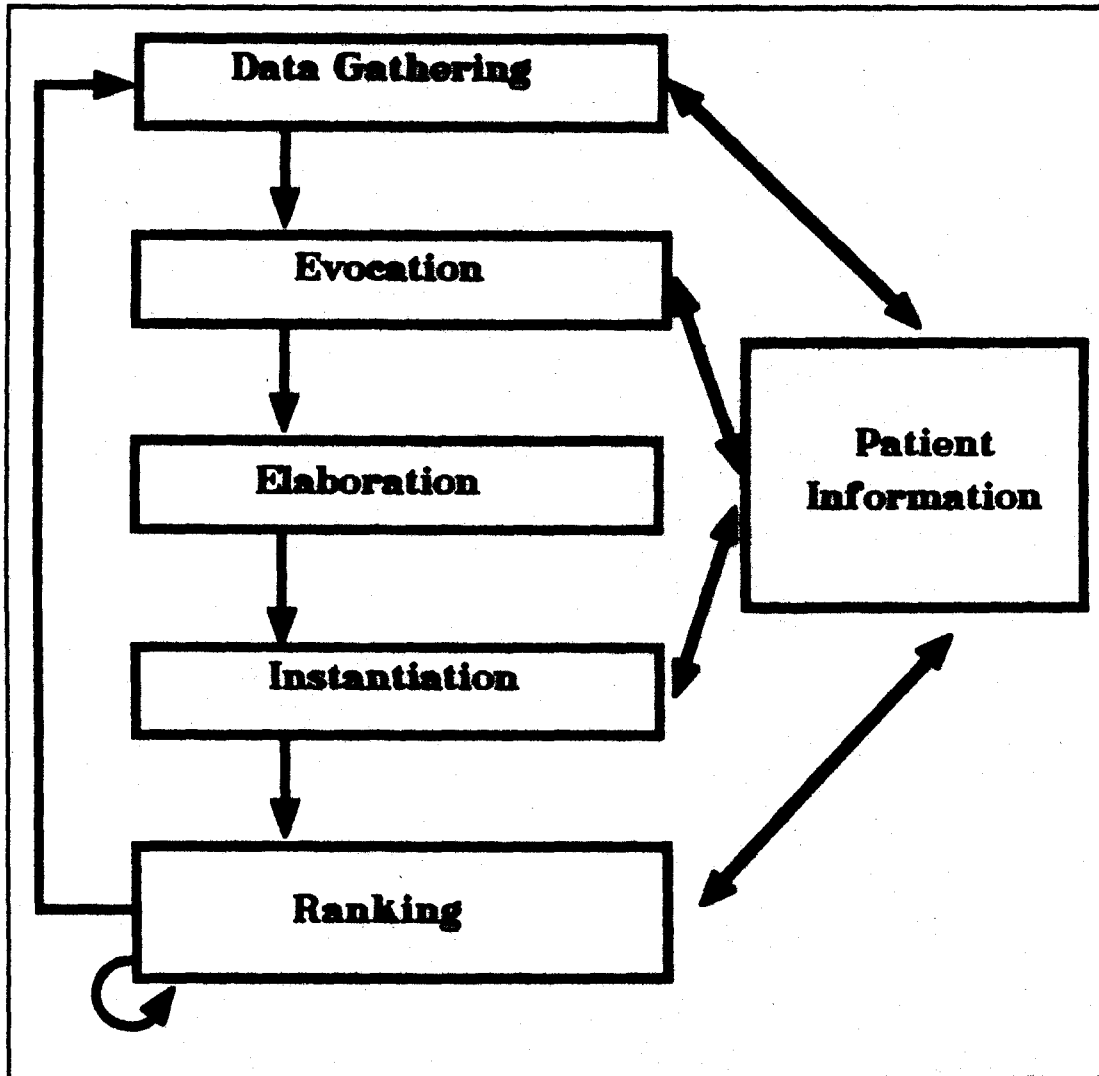


Figure 1.2: Phases in the Taking of a Patient History.

when the line of questioning is directed by the content of the actively considered hypotheses.² The "compiled" data-driven protocol is an evolutionary product of the collective clinical experience which suggests those questions that are always worthwhile asking in response to a particular clinical presentation. For instance, a patient describing chest pain, will usually be asked about its location, radiation, and quality. The temporal data will include the duration of the pain, time of day of onset, temporal relation to precipitating events, and if recurrent, the length of time between episodes.

In general, to avoid being misled by erroneous data, the data-gathering phase, taking place as it does before much investment of diagnostic effort, is an opportune occasion to detect errors. Temporal reasoning can serve in this effort by identifying at least one class of inconsistency. For example, in our earlier transcript of the fictitious office visit, Mrs. W.'s statement that J. was tachypneic two days after he was "fussy" (Sunday) contradicts the assertion that the onset of J.'s anorexia was on Monday and yet occurred at the same time as the onset of tachypnea (Tuesday).

Evocation

Following history-taking, the next diagnostic phase is that of hypothesis evocation or triggering. From the rather large problem-space of diagnostic categories, an expert-system (and for that matter, a human expert) must select a manageably small number of hypotheses.³ This can be achieved by selecting classes of hypotheses based on characteristic constellations of findings in the patient history. For instance, if a patient describes sweating, and a sensation "like a ton of bricks on my chest," this should immediately trigger (bring into consideration) the hypothesis of myocardial infarction (MI). These constellations of findings include characteristic temporal patterns. Chest pain lasting not more than one or two seconds or more than one or two hours is consistent with several diagnostic possibilities, including cervical osteoarthritis, but is atypical for angina.

²This is illustrated in figure 1.2, where the control arrows illustrate requests for patient information at several stages, as well as loops from hypothesis investigation to protocol-driven data-collection.

³See Newell and Simon [42] for a discussion of the limits of cognition and computational resource alike in managing multiple hypotheses.

Elaboration

Once a hypothesis is triggered, the next phase of diagnosis begins, that of elaboration. Each triggered hypothesis usually consists of a broad category of distinct diseases, usually with some pathophysiological pathways in common. Each distinct pathophysiological pathway within the triggered hypothesis is a hypothesis unto itself. The latter hypothesis will be referred to as a *subhypothesis*. The process of extracting the different subhypotheses within the triggered hypothesis is the process of elaboration.

As each subhypothesis represents an alternate clinical course, it is usually the case that each of the subhypotheses has different temporal constraints. A patient with angina pectoris, for instance, may subsequently suffer a sudden MI, or first have congestive heart failure (CHF) before the MI, or immediately suffer a fatal arrhythmia. The duration of the interval from onset of angina to a locally stable state (death or recovery) therefore depends on which clinical course is followed. To return to the office visit example, the duration of J.'s asthma will depend on the subpopulation of asthmatics to which he belongs.

Instantiation

Each of the subhypotheses extracted from the triggered hypothesis is a template for a possible clinical course. To model a particular patient, the data obtained from the patient must be fitted in some way to the subhypotheses, to provide a coherent support for the described condition of the patient. The process of constructing a patient model, matching findings to events in the subhypothesis, is that of instantiation, the fourth of the five phases of patient history based diagnosis.

Patient-specific data further constrains the temporal relations in the subhypotheses, and also anchors the subhypotheses to the present. In J.'s visit to Dr. C.'s office, the patient data specifies that the "current" visit occurred *after* these findings of anorexia and tachypnea were noted, and *before* therapy was instituted. This specification anchors the subhypothesis with respect to the present. This may seem obvious, and yet it prevents unreasonable diagnostic behavior such as seeking therapeutic effect before therapeutic intervention or asking whether J. suffers of asthma as an adult.

Ranking Patient Models

After the creation of several patient models, it becomes possible to proceed to the next diagnostic phase, that of selecting the correct diagnosis or at least ranking these models with respect to their likelihood. Many diagnostic strategies can be brought to bear in this phase to facilitate ranking or confirming a patient model. The goal that these strategies have in common is to determine which datum would be most effective in narrowing the number of diagnostic possibilities or further differentiating the more likely from the less likely.

Temporal information often costs little to acquire and can also serve the various diagnostic strategies⁴ to distinguish between patient models. In distinguishing between an acute and a recent MI, it would be useful to obtain serial SGOT measurements to determine whether the SGOT peak was approaching (the SGOT peak occurs within 24-48 hours of an MI) or past. In the same vein, during J.'s visit to Dr. C.'s office, a history of recurrence, especially in association with a particular season, would support the diagnosis of asthma rather than bronchiolitis. Otherwise, in the acute phase, without these temporal clues, these two clinical entities are often indistinguishable.

1.2 Technical Goals

What then must be required of a temporal representation and reasoner to enable the emulation of the breadth of expression and the behavior described above?

1.2.1 Breadth of Representation

Several efforts in temporal reasoning have made use of heterogeneous internal representations [25,41] to represent the heterogeneity of temporal expression found in the patient history. One of the shortcomings of these representations is the difficulty in determining the order and combination in which the different temporal representations are to be used to retrieve temporal information. The difficulty lies in deciding which combination will be the most efficient, ensuring that the temporal information retrieved is consistent with the rest of the temporal data base and as

⁴The general form of these strategies (e.g. CONFIRM, RULE-OUT, and DIFFERENTIATE) has been described by Patil [45] and Pople [50].

precise as possible. Making this decision becomes far easier if all temporal assertions are maintained within a single representation and manipulated by a single, simple inference mechanism. Consequently, one of the earliest objectives of this research was the development of a uniform underlying representation to see if it could support the necessary breadth of expression.

1.2.2 Computational Feasibility

The computational burden of temporal "reasoning" increases rapidly with the number of events represented in the temporal data base. Several schemes have been developed to work around this problem. Most of the approaches have involved variations on the "divide and conquer" theme. In such cases, the temporal data base is clustered according to a particular scheme, into blocks—usually of contiguous intervals. Temporal reasoning is then restricted locally to each block. Allen's [1] clustering scheme generates hierarchies of intervals where each interval is DURING the interval that is immediately superior to it in the hierarchy. Vere's [58] hierarchies consist of events associated with activities of a planner. The temporal hierarchy built from the planner's activities mirrors the goal-subgoal hierarchy of the planner. In both systems, retrieval of temporal information relating events in different blocks involves some form of search.

Each clustering scheme must be judged by several criteria. First, to what extent do the various application domains fall naturally into the imposed hierarchy? Second, to what degree does the clustering scheme lead to retrieval of inconsistent and imprecise temporal relations? Finally, can the clusters be generated automatically? For pragmatic reasons, a lot of the work on TUP was done to answer the preceding questions.

1.2.3 Interaction of Temporal and Atemporal Reasoning: Domain Independence

It was not at all obvious, at the onset of this project, that temporal data and inferences could be neatly separated from the other conceptual elements of an expert system. Causality, for instance, would appear to intrinsically impose some degree of temporal ordering on events. It was therefore unclear whether temporal reasoning could be isolated and packaged separately or would necessarily have to be tightly coupled with other reasoning mechanisms of a host expert system. If temporal

reasoning could not be isolated in an autonomous package with a standard interface, it would be unlikely that a general solution could be attained to satisfy various expert systems whose representation of causality was as much at variance as say qualitative simulation [29,13,17] and causal association [51,50,46,31]. Achieving a domain and system-independent solution therefore became a major focus of this work.

1.3 Results

Which of the preceding objectives have been achieved? First, the various forms used to specify temporal position within a patient history (heretofore called *temporal descriptors*) have been identified. A uniform representation, based on the *range relation* has been developed, capable of representing the temporal descriptors necessary for producing a patient history. These include quantitative and qualitative relations between any arbitrary combination of points and intervals, references to common temporal yardsticks, references to the present, persistence and alternate temporal hypotheses. A *Temporal Utility Package (TUP)* has been constructed that recognizes assertions of all the patient history temporal descriptors and translates them into range relation form. TUP then performs all the required temporal reasoning and consistency checking, using a form of constraint propagation. TUP also possesses a number of temporal retrieval functions and temporal predicates.

TUP features a standard interface that permits the host system to dictate the desired clustering of the temporal data base for the reasons of computational feasibility mentioned previously. During TUP's development, it became clear that although TUP permits any clustering scheme to be implemented, only the few that conform to a particular set of requirements (discussed at length in the following chapter) will be successful in yielding the desired reduction in computational load. One of these requirements is that events within the same cluster be more closely⁵ related to each other than to events in other clusters. Several expert system technologies, by their nature, provide an organising principle with which temporal clustering can be guided. In this report, an implemented expert system is described that uses its causal aggregation hierarchy (similar to those found in other second generation expert systems such as ABEL [46]) to control the automatic generation of temporal clusters. Planners, frame-based systems, and qualitative simulators also possess

⁵Again, discussion of the exact meaning of this is deferred to the next chapter (section 2.4).

structures that can be employed to automatically guide temporal clustering. These latter instances are discussed in this report but have not been implemented. Also, a number of intriguing cognitive analogies to the temporal clustering schemes appear in the literature. These are discussed in chapter 5.

Reasoning about alternate hypotheses, medical or otherwise, involves examining different outcomes. Each of these outcomes—in medicine, the pathophysiological history—has associated event durations that usually vary with each hypothesis. TUP provides a context mechanism that permits the assignment of temporal assertions to alternate temporal hypotheses. THRIPTH (*Temporal Hypothesis Reasoning In Patient History-Taking*) the medical expert system used as an investigative vehicle, automatically creates contexts for each patient model generated.

1.3.1 Expert System Demonstration

THRIPTH was constructed both to demonstrate TUP's capabilities and to provide a vehicle for investigating the properties of the diagnostic process when the patient models are equipped with a systematic representation of temporal relationships. What THRIPTH really does is best described by returning to the five phases of the diagnostic task, outlined in section 1.1.2. It is in the course of the data-gathering phase, that patient's assertions are processed by TUP and converted into *range relation* form. By virtue of the constraint propagation process that follows each temporal assertion, temporal inconsistencies or contradictions are identified immediately, and the user/patient given the option of which assertion to withdraw.

Hypotheses are evoked in the following phase by means of "triggers" linked to hypothesis templates. Each trigger tests the patient assertions accumulated in the data-gathering phase. Naturally, the triggers can test for temporal relations that are deduced from the assertions, even if not explicitly asserted. The triggers themselves are arbitrary boolean combinations of temporal and atemporal predicates.

Each hypothesis template is represented as a directed, acyclic graph of event nodes linked by associational or causal links. Within each such template, exclusion relations specify which events are mutually exclusive. These exclusion relations are used in the hypothesis elaboration process to generate distinct and mutually exclusive subhypotheses.

Temporal information is associated with each event in the hypothesis template. This information includes specification of the temporal antecedence of causally antecedent events, event durations and other *a priori* temporal knowledge. On elab-

oration, the temporal information is carried forward to each subhypothesis and a temporal context generated from this information. It is at this point, during the generation of temporal contexts, that the temporal data is clustered. THRIPTH communicates to TUP the causal aggregation information that TUP then uses for defining the temporal clusters.

During instantiation, the patient data is bound to each subhypothesis. This involves employing the temporal relations, obtained in the initial history-taking, to modify the default assumptions of the subhypotheses.

As outlined previously, the primary goal of the diagnostic strategies of the fifth phase of patient history-driven diagnosis is to find features of the competing hypotheses that distinguish them from one another and then to obtain patient data related to the differing expectations of the respective patient models. Timing, of course, is one such feature and therefore in addition to distinguishing patient models by the presence or absence of supporting findings, the diagnostic strategies seek differences in temporal distances between events. This diagnostic strategy loop has only been partially implemented, principally because of the effort required to elaborate such strategies, even without temporally sophisticated capabilities.

1.4 Organization

The following chapter describes the representation and reasoning mechanism of TUP. Comparisons to previous work in this area are made as the issues arise. I also discuss, in general terms, how TUP would interact with various expert system technologies.

Chapter 3 describes the specifics of a temporally oriented expert system—THRIPTH—and the temporally explicit patient record, and how these interact with TUP.

In addition to the benefits of added expressive power and generality of reasoning that temporal reasoning can bring to an expert system, developing an expert system, so endowed, with domain knowledge, places a much larger burden on the knowledge engineer. In chapter 4, the challenge of extracting temporal knowledge from the literature and delivering it to an expert system is examined.

There have been several studies, in the discipline of cognitive science, of temporal reasoning. It becomes tempting, because of similarities that appear, to draw analogies between the experimental results of cognitive science and the behavior of TUP. However, the results are not without controversy, even within the ranks

of cognitive scientists, as is the relevance of these experiments to human behavior. More on this in chapter 5.

Chapter 6 recapitulates the results of the work. Areas of weakness are noted and promising directions for future work in this area are pointed out.

Appendix A is a compilation of several tables, too large to include in the main body of text, and includes conversions between point and interval definitions, and the definitions of TUP predicates, assertions and retrieval functions. Appendix B contains a list of temporal assertions used in a causal hypothesis. Appendix C provides selected definitions of the abbreviations used in this report.

2. The Temporal Utility Package

The *Temporal Utility Package* (TUP) is designed to perform domain and system-independent temporal reasoning. TUP's internal representation and reasoning mechanisms are described in this chapter.

2.1 The Range Relation

The *range relation* (or RREL in TUPese) is a first-order object in TUP's representation. It is from this object that all the other types of temporal representation are built. Range relations specify the upper and lower bounds on the temporal distance between two points in time. The range relation is specified in TUP's syntax using the form below:

Form 1

```
(RREL <first point> <second point>  
      <lower bound> <upper bound>  
      <context>)
```

Individual points are identified in this form by lists of qualifier-value pairs. The qualifiers can specify whether the point represents a point event or the beginning or end of an interval. Qualifiers may also specify that a point is a member of a cluster (see section 2.4) or is a point along a particular temporal yardstick (see section 2.6).

Bounds on the range relation can have all values from positive to negative infinity. Numerical values can be in any time unit from seconds, minutes, to centuries, even though all these values are canonically stored in seconds. The value of the lower bound must however never exceed that of the upper bound. If TUP makes a temporal deduction that violates this rule (see section 2.2.4), it is always taken to mean that there are inconsistent or contradictory temporal assertions.

As mentioned in the introduction, contexts permit the representation of alternate temporal assertions. The context specified in the RREL assertion of form 1 can be any object—atomic symbol or hypothesis. If not specified in an assertion, it defaults to REALITY.

2.1.1 Diagrams and Descriptions

Some explanation is required for the diagrams and English "translations" that accompany the temporal assertions in the examples that follow. The diagrams will initially only serve to illustrate the graphic format used to describe the RREL. However, as the assertions increase in number and more complex features are demonstrated, these diagrams will become the primary vehicle for analysing the consequences of the assertions. The mapping from RREL to English is a delicate matter and is the major reason for the graphic illustration of the temporal relationship. The problem stems from the richness of meaning of the English language. This richness makes it all too easy to attach meaning to temporal assertions other than is present in the formal syntax (or diagram). In example 1 (page 19) for instance, I am careful to avoid making any allusion to the position of irritability or anorexia with respect to the present.¹ Specifying the position with respect to the present would require the assertion of at least one more RREL (as described in section 2.1.2). Consequently, the English translation provided should be viewed as an aid in comprehending the temporal relations; the graphic representation however should be examined to obtain the strict interpretation.

Regarding the diagrams themselves, two different graphical representations of the temporal relations were considered—time-line and directed graph, both illustrated in Figure 2.1. In time-line diagrams the points are visually ordered with respect to a reference point, with a scaled lower and upper bound displayed for each point. The directed graph diagrams do not necessarily visually convey information of the temporal order, but rather permit inspection of individual point to point temporal relations. Time-line diagrams have several weaknesses, at least for the purposes of illustrating reasoning with RRELs. As all temporal relations are illustrated with respect to a reference point, it is difficult (without additional diagrams that use other reference points) to appreciate temporal relations other than those that involve the reference point. This is especially true when the upper and lower bounds do not share the same sign.² Also, to understand how TUP makes its temporal deductions, it is important to visualise the connectivity between the RRELs. This is particularly useful in following how TUP clusters the temporal data base.

For the reasons above, except for example 1, I have illustrated temporal relations

¹As in "preceded" or "will precede."

²i.e. a point occurs before or after the reference point.

exclusively with the directed graph format. However, to the degree it is possible, I have attempted to visually convey the ordering information using the directed graph format, by having earlier points above, or to the left of, later time points.

2.1.2 Examples

The description of TUP's reasoning mechanisms will be a lot clearer and better motivated if we first look at several RREL examples.

Example 1

```
(RREL ((NAME IRRITABILITY)(TYPE BEGIN-INTERVAL))
      ((NAME ANOREXIA)(TYPE END-INTERVAL))
      (2 DAYS) (3 DAYS))
```

Example 1 corresponds to the assertion that the beginning of irritability precedes the end of anorexia by two to three days.

Example 2

```
(RREL ((NAME SLEEP-TERRORS)(TYPE BEGIN-INTERVAL))
      ((NAME AWAKE)(TYPE END-INTERVAL))
      (-3 HOURS)(-1 HOURS))
```

Example 2 corresponds to: the onset of sleep terrors occurs one to three hours after the end of awakeness.

Example 3

```
(RREL ((NAME ASTHMA-1) (TYPE END-INTERVAL))
      ((NAME ASTHMA-2) (TYPE BEGIN-INTERVAL))
      (+EPSILON) (+INFINITY))
```

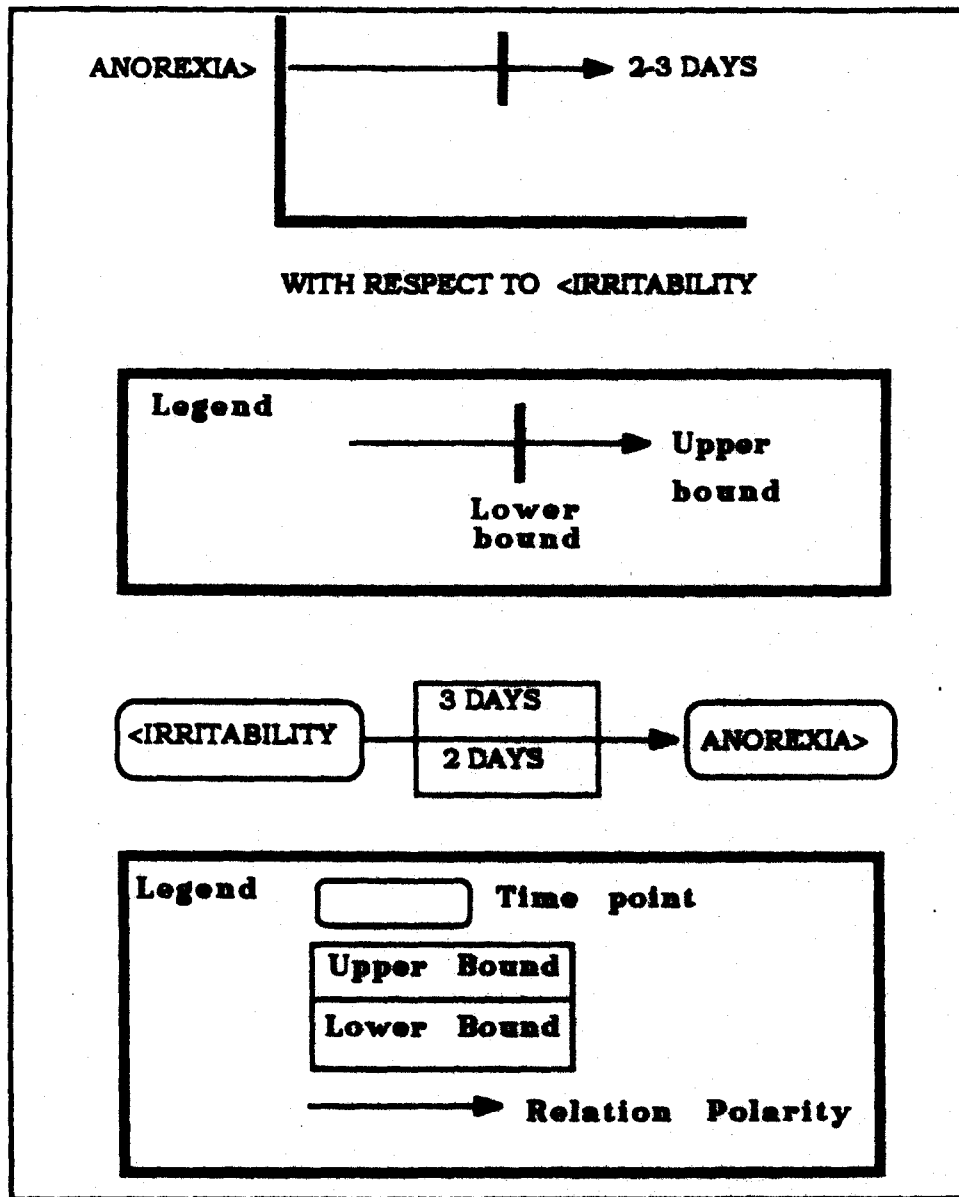


Figure 2.1: Time-Line and Graph Form of Temporal Diagram

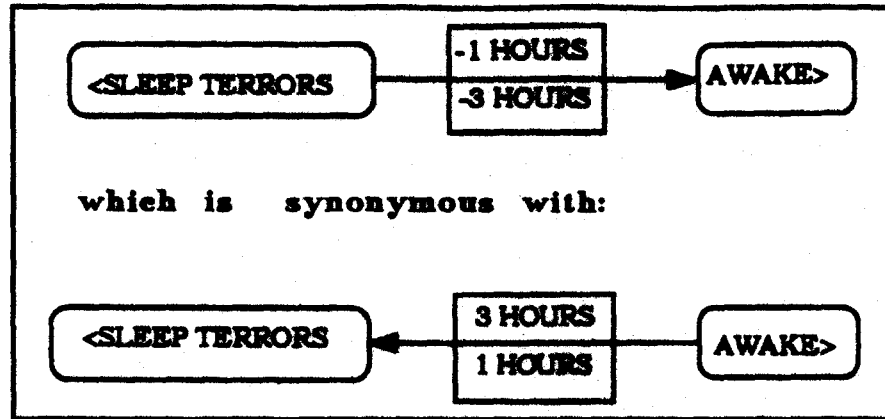


Figure 2.2: Illustration of assertion of example 2.




Figure 2.3: Illustration of assertion of example 3.

This assertion has a lower bound ϵ , a quantity that TUP takes to be smaller than any numerical value the computer can represent. Temporal reasoning with such values is described in section 2.2.2. The assertion of example 3 corresponds to: the attack of asthma episode number two can occur any time after³ asthma episode number one. Note that although the upper bound of the RREL is $+\infty$, in any complete expert system that would use TUP, this bound would be constrained to the maximum life expectancy of the individual. In THIRPHT's knowledge bases, many event intervals are asserted to occur during the interval of LIFE. Since LIFE (really the RREL that relates the beginning and end of the interval) has an upper bound of longevity generously set at 120 years and 237 days [52], most temporal relations of pathophysiological events become finite after constraint propagation (described in section 2.2.2).

Example 4

```
(RREL ((REFSYS CALENDAR)
      (REFSYSFORM "8 p.m., Sunday, April 27th, 1986")
      (TYPE POINT))
      ((NAME IRRITABILITY)(TYPE BEGIN-INTERVAL))
      (-1 HOURS)(+1 HOURS))
```

Example 4 introduces the temporal "yardstick". Although any number of such yardsticks can be defined, at the moment TUP only includes definitions of the calendar, ages, developmental stages, and life-landmarks. Within TUP these yardsticks are implemented as mini-experts that know enough about their own reference system (the temporal yardstick) to be able to assert RRELS between points belonging to the same reference system—discussion of details is deferred until section 2.6. The CALENDAR reference system recognizes all forms that evaluate to points on the date-line whereas DEVELOPMENT accepts all defined stages of development including embryonic, fetal, infancy, early and late childhood, early and late adolescence and adulthood. In referring to a temporal yardstick, TUP requires a form (REFSYSFORM) that evaluates to a point on the yardstick and a selection of yardstick (reference system or REFSYS in TUPese). Example 4 can therefore be interpreted as: the onset of irritability occurs between 7 and 9 p.m. on April 27th, 1986.

³From immediately after, to infinitely far in the future.

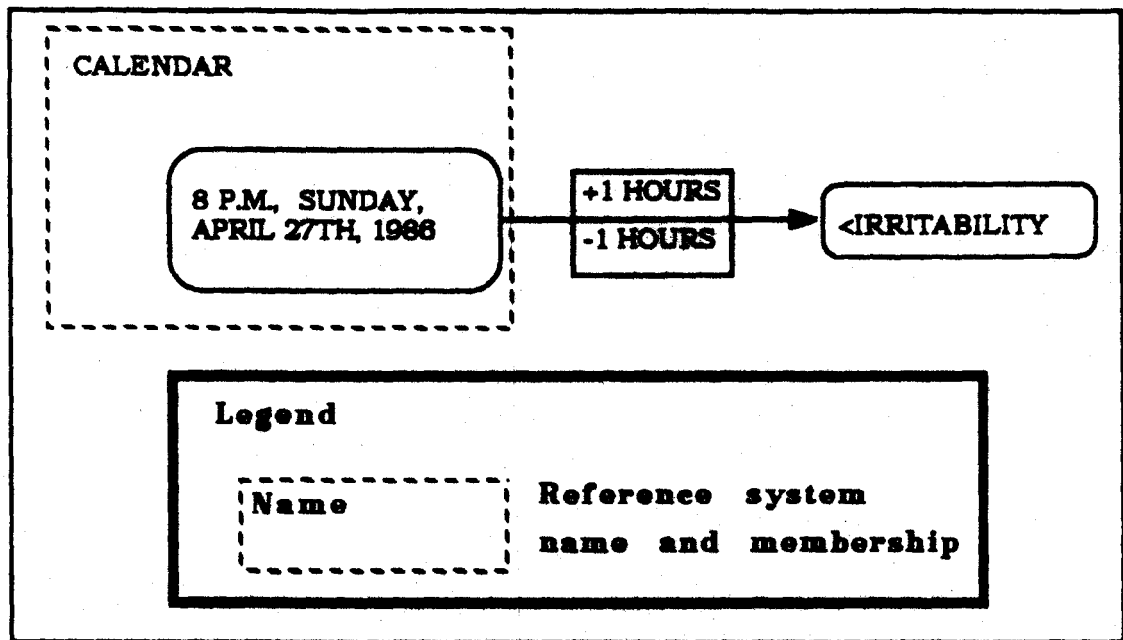


Figure 2.4: Illustration of assertion of example 4.

Example 5

```
(RREL ((NAME ASTHMA-PRONE)(TYPE END-INTERVAL)
      ((REFSYSFORM "END-ADOLESCENCE")
       (REFSYS DEVELOPMENT) (TYPE POINT))
      (0 YEARS) (0 YEARS)
      CONTEXT-1)

(RREL ((NAME ASTHMA-PRONE)(TYPE END-INTERVAL)
      ((REFSYSFORM "END-ADOLESCENCE")
       (REFSYS DEVELOPMENT) (TYPE POINT))
      (-INFINITY) (-EPSILON)
      CONTEXT-2)
```

An extremely simple use of the context mechanism is illustrated by example 5. Two alternate assertions are given to describe the relation between the end of adolescence and the end of the asthma-prone interval. In the first assertion, in **CONTEXT-1**, the asthma-prone period ends with adolescence; in the second (**CONTEXT-2**) it ends any time after the end of adolescence. Whereas in example 5, the context slot is bound to an isolated atomic symbol, it is used by **THRIPHT** to bind full-fledged, structured hypotheses.

Although the role of asserting temporal position with respect to the present will not become apparent until the section dealing with temporal reasoning, example 6 presents a preliminary introduction. The RelationToPresent assertional form (form 2) is translated to the two RRELs shown. The first RREL establishes the relation between the current instance of the present (represented by a point whose event name is a unique NOW) and the onset of asthma. The second RREL asserts the relation between this latest instance of the present and the current time obtained from the host computer. The latter point is treated identically to other points on the calendar reference system. Example 6 then roughly corresponds to: asthma began twelve to thirteen months ago.

Form 2

```
(RelationToPresent
  <point> <lower bound> <upper bound> <context>)
```

Example 6

```
(RelationToPresent
  ((NAME ASTHMA) (TYPE BEGIN-INTERVAL))
  (+12 MONTHS) (+13 MONTHS))
```

translates to:

```
(RREL ((NAME (GENSYM NOW)) (TYPE POINT))
  ((NAME ASTHMA) (TYPE BEGIN-INTERVAL))
  (-13 MONTHS) (-12 MONTHS))

(RREL ((REFSYS CALENDAR)
  (REFSYSFORM (DATE)) (TYPE POINT))
  ((NAME (LatestNow)) (TYPE BEGIN-INTERVAL))
  (0 SECONDS) (0 SECONDS))
```

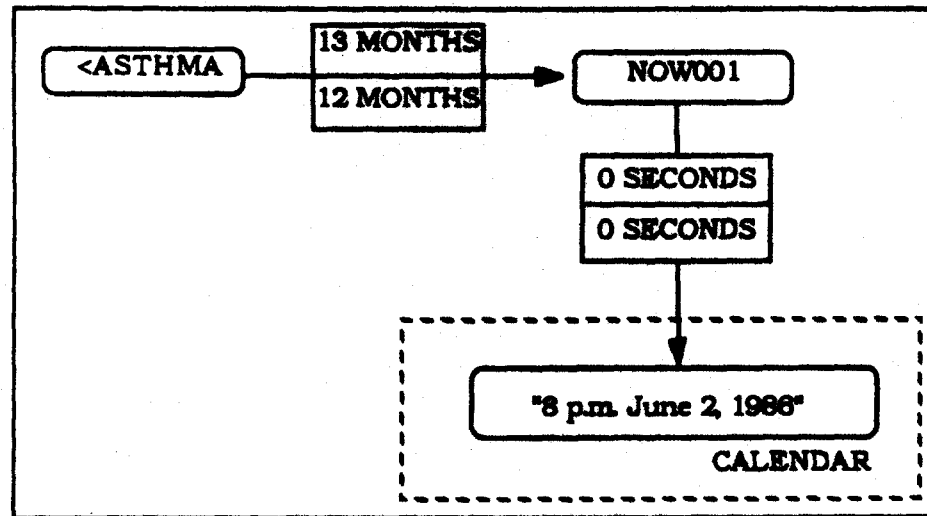


Figure 2.5: Illustration of assertion of example 6.

On many occasions, one may wish to directly assert temporal relations between intervals without explicitly specifying the position of the onset and end of each interval. Just such a functionality is provided by the interval-based temporal reasoners, developed by Allen [1] and Vilain [61], that employ thirteen different interval relations. These thirteen relations can be conveniently used in TUP (form 3) as shown in example 7.

Form 3

```
(INTREL
  <interval one> <interval two>
  <interval relation>
  <context>)
```

Example 7

```
(INTREL ((NAME IRRITABILITY)) ((NAME ANOREXIA)) OVERLAPS)
```

translates to:

```

(RREL ((NAME IRRITABILITY) (TYPE BEGIN-INTERVAL))
      ((NAME ANOREXIA) (TYPE BEGIN-INTERVAL))
      (+EPSILON) (+INFINITY))

(RREL ((NAME IRRITABILITY) (TYPE END-INTERVAL))
      ((NAME ANOREXIA) (TYPE END-INTERVAL))
      (+EPSILON) (+INFINITY))

(RREL ((NAME IRRITABILITY) (TYPE END-INTERVAL))
      ((NAME ANOREXIA) (TYPE BEGIN-INTERVAL))
      (-INFINITY) (-EPSILON))

```

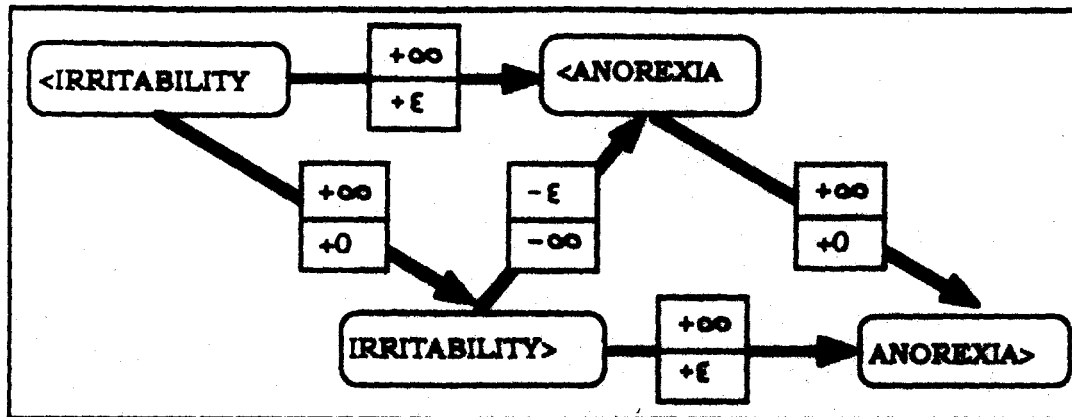


Figure 2.6: Illustration of assertion of example 7.

Note that with the INTREL assertion, TUP automatically asserts the additional default temporal relations between the beginning and end of each interval (using ASSERT-INTERVAL as in form 4).

All thirteen conversions from interval to point-based representation are tabulated in appendix A. TUP also provides a short-hand means of creating intervals, the ASSERT-INTERVAL form (form 4).

Form 4

```
(ASSERT-INTERVAL <interval specification> <context>)
```

translates to:

```
(RREL ((<point specification>) (TYPE BEGIN-INTERVAL))
      ((<point specification>) (TYPE END-INTERVAL))
      (0 SECONDS) (+INFINITY)
      <context>)
```

The examples above show a variety of temporal assertions that are represented internally in a single uniform manner. The method by which this internal representation performs temporal reasoning is discussed in the following section.

2.2 Generating Temporal Inferences

Every time additional temporal information is acquired, new relations may be generated or any of the temporal relations, currently represented in the data base, modified. In this section, I describe how these inferences are computed.

2.2.1 Range Addition

If an RREL is asserted that shares a point with another RREL, TUP attempts *range addition* to calculate the bounds on a third RREL—that between the two unshared points. Range addition is simply the calculation of the sum of the two lower bounds and the sum of the two upper bounds. The rules for range addition for all values including ∞ and ϵ are given in table A.2 in appendix A.

Let us take, for instance, the assertions of example 8, in which the assertion is made that the onset of irritability precedes the onset of anorexia by two to three days and that the duration of anorexia is three to four days. The lower and upper bounds on the third relation, (diagrammed in Figure 2.7), as calculated by range addition, are five days and seven days respectively.

Example 8

```
(RREL ((NAME IRRITABILITY) (TYPE BEGIN-INTERVAL))
      ((NAME ANOREXIA) (TYPE BEGIN-INTERVAL))
      (2 DAYS) (3 DAYS))

(RREL ((NAME ANOREXIA) (TYPE BEGIN-INTERVAL))
      ((NAME ANOREXIA) (TYPE END-INTERVAL))
      (3 DAYS) (4 DAYS))
```




Figure 2.7: Temporal Data Base After Design Addition.

2.2.2 Constraint Propagation

If an RREL r_2 , such that $i \xrightarrow{r_2} j$, is asserted between two points i, j which are already related by $i \xrightarrow{r_1} j$ TUP determines if r_2 constrains r_1 . An RREL is constrained if the values of its upper and lower bounds are brought closer together, be it by increasing the lower bound, decreasing the upper bound or both. A state of maximum constraint is achieved when the values of the upper and lower bounds are identical. If the lower bound is greater than the upper bound, the temporal data base is in internal contradiction (is inconsistent).

To continue with the example, if it is subsequently asserted that the delay from the onset of irritability to the end of anorexia is exactly five days (as in example 9), then the RREL originally computed can be modified—or constrained. Here, the bounds on the new assertion (example 9) are narrower than those of the RREL that was obtained through the range-addition of the RRELS of example 8. And so, the RREL is constrained (see Figure 2.8).

Example 9

```
(RREL ((NAME IRRITABILITY) (TYPE BEGIN-INTERVAL))
      ((NAME ANOREXIA) (TYPE END-INTERVAL))
      (5 DAYS) (5 DAYS))
```

Every time an RREL is constrained, range addition is attempted with each of the RREL's neighbors (RRELS with one time point in common). This range addition may either constrain a previously asserted RREL, create a new one if no prior one exists, or do nothing if the range addition does not constrain extant bounds. In this example, the first neighbor that might be chosen is the first RREL of example 8. This results in a range addition of (-3 days, -2 days) and (5 days, 5 days) for a result of (+2 days, +3 days). Fig. 2.9 shows that this result constrains the upper bound of the RREL between the beginning and end of anorexia to three days. As this process of constraint and range addition is recursively repeated every time an RREL is constrained (whereby constraint propagation), this latest constraint causes yet another range addition: (-3 days, -3 days) and (5 days, 5 days) with a result of (2 days, 2 days) that constrains the corresponding RREL as illustrated in Fig. 2.10.

Although TUP attempts the remaining range additions, there obviously cannot be any additional constraint propagation as all three RRELS are maximally constrained. As Vilain and Kautz [59] have shown, such a constraint propagation scheme has order n^3 time complexity. In outline, constraint propagation causes

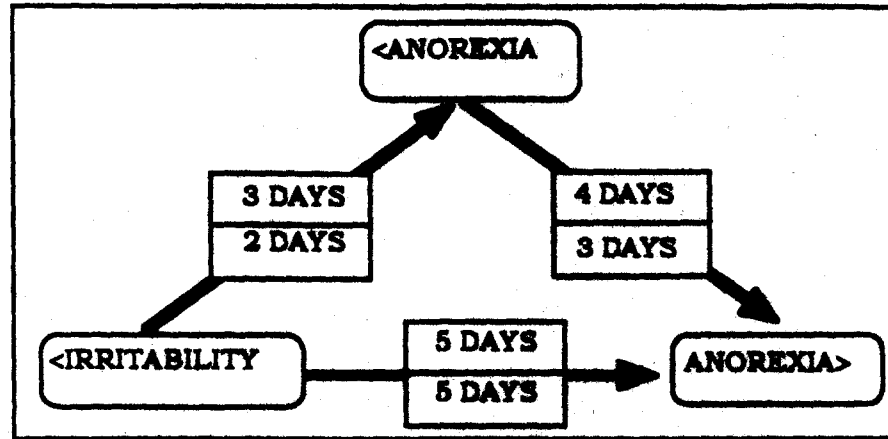


Figure 2.8: Initial Constraint.

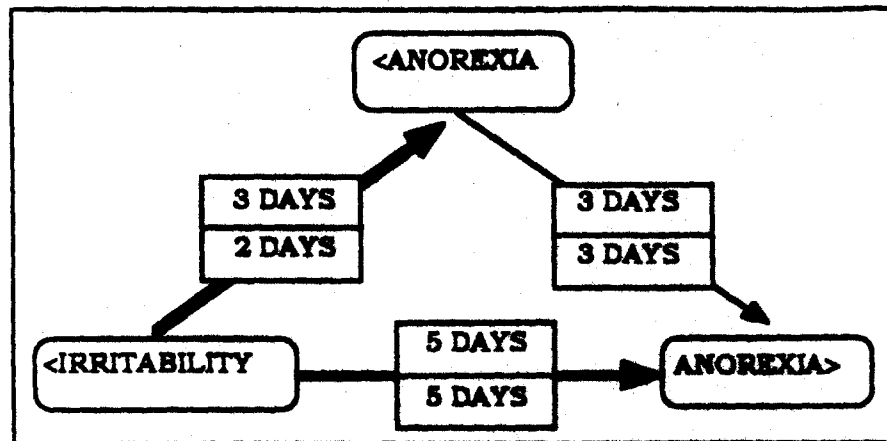


Figure 2.9: First propagation of constraint.




Figure 2.10: Final propagation of constraint.

points to be exhaustively interconnected with RREs. Consequently, every time an additional RRE is added to the data base, there are order n^2 other RREs which might be constrained as a result. To build up each of n RREs n of the n^2 constraint operations may be executed. The problem grows in n^3 complexity.

We can now return to the example of the beginning of this chapter to show the effects of TUP's constraint propagation upon the temporal data base. With constraint propagation, a clock RRE is generated by the RREs of example 7, as illustrated in Fig. 2.11. If we further connect the RREs with the RRE of example 10, the result, shown in Fig. 2.12, is obtained.

Example 10

```
(RRE ((NAME IRREVERSIBILITY (TYPE IRREVERSIBILITY))
      (START 1970-01-01) (END 1970-01-01))
      (NAME IMMUTABILITY (TYPE IMMUTABILITY))
      (START 1970-01-01) (END 1970-01-01))
```

2.2.3 Temporal Dependence

TUP maintains, in its representation of RREs, a record of the range additions used to calculate the bounds for each RRE. Accordingly, in TUP's object representation, each RRE r_i contains a record of the range of generation of the form:

< r_1, r_2, A, \dots >

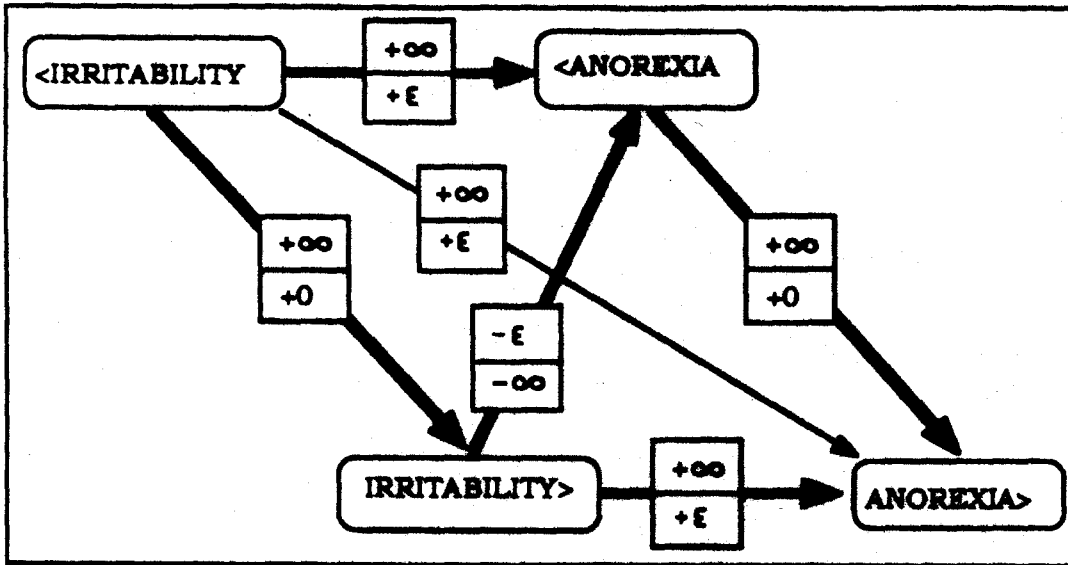


Figure 2.11: Propagation of the "OVERLAPS" Assertions.

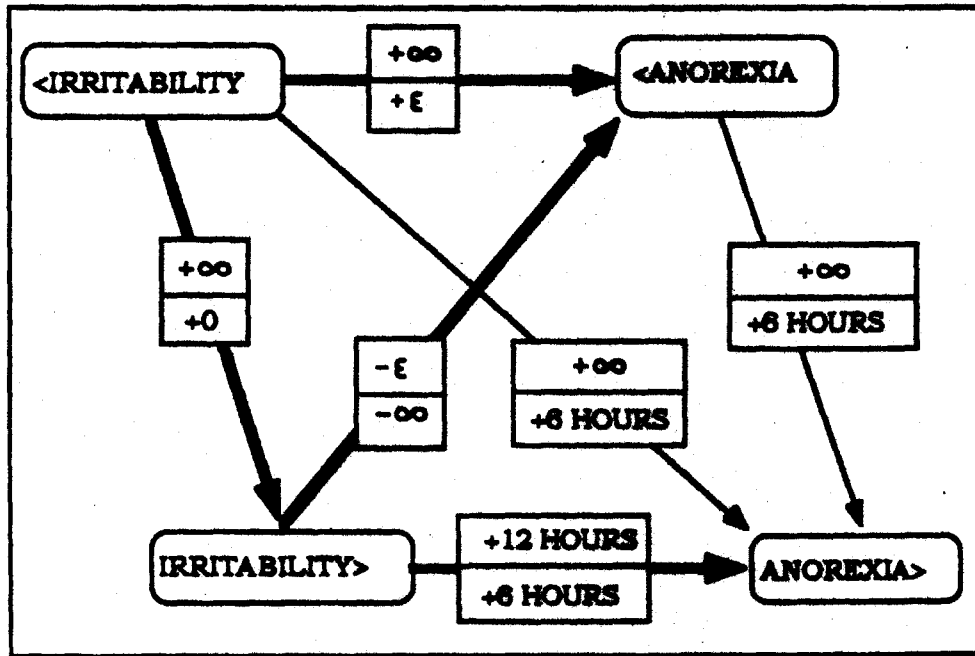


Figure 2.12: Constraining the "OVERLAPS" assertion.

where r_2, r_3 are RRELS which were used in a range addition to calculate the bounds on r_1 , and lb, ub are the results of that range addition. These supports are used, as described in the following section, to trace the inference paths from any RREL to the original external assertions on which the bounds of that RREL depends. A similar scheme was described by McDermott [36]

2.2.4 Contradiction Handling

If, during the process of constraint propagation, a contradiction is detected, the contradiction handler is invoked. The contradiction handler's basic strategy is to offer a list of externally asserted RRELS for withdrawal and to continue to do so until no more contradictions are detected.

Withdrawal of an assertion involves more than merely deleting a single RREL. TUP has to examine all other RRELS for which the externally withdrawn⁴ RREL (that we will call the *inconsistent RREL*) was a support. If this withdrawal of support causes a decrease in constraint in a second RREL, then this RREL must in turn be (temporarily) withdrawn from all supports. This process recurses until there ceases to be any further decrease in constraint. Subsequently, all those RRELS whose constraint was weakened by the withdrawal of the inconsistent RREL are automatically re-asserted by TUP (with consequent constraint propagation).

This process of dependency-directed backtracking is computationally expensive, and is therefore only intended for use in small applications involving only a few hypotheses. Otherwise, for large, realistically-sized applications, such as medical diagnosis, the context mechanism is used as it permits several hypotheses to be repeatedly examined and compared with each other at minimal cost.

One of the more challenging tasks, as yet unresolved, is the automatic selection of the inconsistent RREL to be withdrawn. One method is to attempt to withdraw those premises, of those in mutual contradiction, that are least supported by other RRELS. As many such methods that solely rely on the domain-independent, syntactic information of the temporal data base, this method is incomplete and unreliable. Domain knowledge is required to understand why a particular temporal assertion is unreasonable. We know for instance, that acute hepatitis B would have to follow rather than precede infection with the Hepatitis B virus, but this would not be obvious from the temporal relations alone—we had to know something about hep-

⁴That is, withdrawn explicitly by user or expert system rather than by TUP's inference mechanisms.

atitis, or infections in general. It is precisely to provide this external or "real-world" knowledge that THRIPT is equipped with temporal templates for each hypothesis (described in section 3.4).

Linear Programming and Contradiction Handling

Linear programming has been considered as an alternative to constraint propagation [34] to perform temporal deduction. As noted by Valdés-Péres [57], full linear programming presents some difficulties. One of these is the problem of defining which sets of assertions are in contradiction. TUP determines this by tracing back the dependencies that lead to the contradiction; it is not apparent how this can be accomplished using the Simplex algorithm. Other drawbacks of the full linear programming techniques include what Valdés-Péres terms the "lack of implementational congeniality." That is, if we represent events using some particular representation, then the use of Simplex requires some properties of these events to be represented in a set of wholly different data structures. Integrating the two representations can be at the very least difficult, and leads to knowledge structures that are less than obvious to the knowledge engineer. Also, the complexity and unintuitive nature of the operations of the Simplex algorithm does not permit an expert system which uses the algorithm to readily generate an explanation of its deductive process. This contrasts with the deductions made through constraint propagation, where a backwards trace through the supports of an RREL provides a reasonable explanation of how the bounds were computed.

2.2.5 Tools for Studying Constraint Propagation

During TUP's design, I developed some tools to follow the consequences of constraint propagation upon the temporal relations already entered in the knowledge base. I have subsequently found these tools to be useful in guiding temporal knowledge engineering (see chapter 4).

Loop Formation

A loop in TUP is constituted of a set of n externally asserted RRELS joining n time points to form a closed graph as illustrated in Figure 2.13. The simplest loop is that with three RRELS, but of course it can involve any number of RRELS. The interest of loop formation is that it is only when loops are formed that RRELS are

constrained⁵ and therefore that the externally asserted RRELS are modified. If this is not immediately apparent, consider the simplest case—two RRELS, as in example 8. The only circumstance in which a third, distinct, externally asserted RREL will cause constraint propagation is if it joins the two points (BEGIN IRRITABILITY) and (END ANOREXIA) to form a loop. This third RREL, which closes the loop, provides an alternate path for the computation of the temporal distance between the two points.

In the general case, when an externally asserted RREL closes a loop of any size, it provides an alternate computation path for the bounds of all the other RRELS that are members of the loop. Thus in the simplest case, the closure of the loop with the third RREL can constrain the two previously asserted RRELS. Constraint propagation will also spread throughout larger loops and can propagate to the members of intersecting loops.

In the current implementation, TUP highlights, on generated graph diagrams, externally asserted RRELS. This enables the visual detection of loops by the knowledge engineer, and permits her to direct her efforts in knowledge acquisition to those RRELS that might close loops and therefore constrain the temporal data base. I have found this particularly helpful in modifying underconstrained disease hypotheses.

The Constraint Index

The constraint index (CI) is a measure of the constraint of an RREL. It is calculated by dividing the distance between the upper and lower bounds of an RREL by the sum of the absolute value of the same bounds.

$$CI = \frac{|ub - lb|}{|ub| + |lb|}$$

For instance, the first RREL of example 8 has a constraint index of (3 - 2) days divided by (3 + 2) days or 0.2. By the formula, the more constrained an RREL is, the lower its constraint index.

The CI is calculated so that even when two RRELS have bounds separated by an identical distance, the RREL that covers a larger time scale has the lower CI. In this way the RREL that relates events years apart is rated as more constrained than one relating events months apart even if the difference between the upper and lower bounds of both is an identical number of days.

⁵Other than asserting an RREL that directly constrains a previously externally asserted RREL.




Figure 2.13: Loop Formation

TUP can calculate the cumulative sum of the CI's for specified groups of points (form 5). This is known as the group constraint index (or global constraint index when it includes the whole temporal data base).

Form 5

(CI <point list>)

The CI is useful for the knowledge engineer, because it enables quick identification of those parts of the temporal knowledge base that are poorly constrained. With the graphical display of loops, it provides an efficient tool for building highly constrained knowledge bases. Also, it can be used for the performance clustering heuristic, described in section 2.5.4.

2.3 RREL Retrieval

If, instead of values between positive and negative infinity, the bounds of an RREL are given as TUP variables as in form 6, TUP attempts to retrieve the values for those bounds instead of asserting them and binds the retrieved values to the variables. If only one of the bounds is an unbound TUP variable, TUP first asserts the RREL replacing the variable bound with infinity (positive if upper bound and negative if lower bound), performs the retrieval operation, and then binds the retrieved value of that bound to the variable.

Form 6

(RREL <point 1> <point 2> ?<LB-var> ?<UB-var> <context>)

In a temporal knowledge base with exhaustive constraint propagation, this retrieval computation simply returns the one RREL linking the two points specified. As described below (section 2.4), the retrieval process becomes somewhat more involved with clustering.

TUP also provides retrieval functions for temporal relations built on the RREL. For instance, if an unbound TUP variable instead of an interval type is specified in an INTREL, TUP tests the RRELs between the two intervals and returns the interval relation that is consistent with the RREL relationships. If, as is often the case, the

RREs are insufficiently constrained, TUP returns a disjunction of interval relations that are consistent with that state of constraint. Conversions from range relations to interval relations can be found in table A.3.

2.4 Restricting Constraint Propagation

Temporal reasoning, as described in the preceding section, is uniform and straightforward. Unfortunately, computation of constraint propagation becomes burdensome in typical medical applications where each hypothesis contains hundreds of events. With the first attempts at developing applications that use constraint-propagating temporal reasoners, more attention has been paid to the tractability issue. It has become apparent that the time complexity is in large part due to the attempt to achieve breadth of expression while ensuring complete and consistent temporal inference. Valdés-Pérez [57] and Vilain and Kautz [59] have examined this tradeoff and have also developed representations that permit inference with order n^3 time complexity while maintaining global consistency.

Unfortunately, even n^3 complexity is much too high. To process, in real-time, several simultaneous hypotheses, each involving hundreds of events, requires a much shallower performance curve. The improved performance is necessary as TUP is meant to be an accessory program for an expert system and not the primary computational activity.

Other than this pragmatic objection to unrestricted constraint propagation, there are several reasons to believe that such constraint propagation is not necessary.

With unrestricted constraint propagation, all temporal relations are used to obtain maximum constraint. Nonetheless, experience with TUP has shown that in practice, only a small fraction of temporal relations in the data base need be used to obtain maximum constraint (more on that later in this section). Perhaps non-coincidentally, this observation dovetails with our intuition that we do not use or review all our remembered events to estimate the temporal distance from one event to another. This intuition appears to be borne out by cognitive experiments described in chapter 5.

Most temporal reasoners use some form of "divide and conquer" to allow temporal inference to be performed on subsets of the events (clusters of the temporal data base) that are small enough to prevent the computational burden from becoming onerous. In those temporal inference engines that use constraint propagation, the

usual approach is to select a cluster of events and then precompute (through constraint propagation) the relations between them. Thus, to overcome the complexity of his interval-based constraint propagation algorithm, Allen [1] uses the *reference interval* to cluster intervals in a *during* hierarchy. Constraint propagation is limited to reference intervals such that relations between intervals in different reference intervals must be obtained by means of a search procedure. Allen notes that "If one is careful about structuring the reference hierarchy, this [local constraint propagation and the search procedure] can be done with little loss of information from the original complete propagation scheme."

Similarly, to manage the steep growth in execution time of the DEVISER I planner, Vere [58] employs a form of temporal clustering. DEVISER II associates a temporal scope with each goal activity. This allows the planner to exploit the goal hierarchy by selectively retrieving events within the temporal scope of the goal activity. One of the advantages of Vere's solution is that it does not require the knowledge engineer to perform the difficult and laborious task of temporal clustering.

Unlike Allen's and Vere's temporal reasoners, those of Mittal [41] and Kahn [25] both have heterogeneous representations and cluster schemes. Mittal's temporal data base maintains three nested types of temporal organization: event clusters (events clustered around a selected time), episodes (event clusters organized around key events) and episode clusters (clusters of recurrent episodes). All three of these aggregation methods perform a role equivalent to clustering; by grouping events according to some particular criteria, the retrieval mechanism can quickly focus on a small subset of the total data base. Kahn's [25] temporal reasoner has three parallel representations: before-after chains, the date-line and the position of events relative to key events. The third representation clusters those events for which the key event has shared importance.

The heterogeneous internal representations suffer from fundamental difficulties in retrieval, namely the problem in deciding which representations to use and in which order to use them to obtain the most precise temporal information.⁶ In their implementations, such data bases use several heuristics to make retrieval decisions, for example employing those representations or representation combinations that are most frequently successful.

In the sections that follow, I describe how TUP implements clusters and then

⁶e.g. it is *a priori* difficult to decide whether the effort of searching before-after chains will provide more precise information than searching the date-line to respond to a particular query.

how these clusters can be automatically generated.

2.4.1 Implementation of Clustering: The Reference Set

In the point specification in an RREL assertion, one of the qualifiers available is the REFSET that specifies reference set membership. Events can belong to more than one reference set, and therefore a reference set qualifier can specify a list of memberships as in example 11.⁷

Example 11

```
(RREL ((NAME DYSPNEA) (TYPE BEGIN-INTERVAL)
      (REFSET (RHEUMATIC-HEART-DISEASE MITRAL-STENOSIS)))
      ((NAME RHEUMATIC-FEVER) (TYPE BEGIN-INTERVAL)
      (REFSET (RHEUMATIC-HEART-DISEASE)))
      (-49 YEARS) (-47 YEARS))

(RREL ((NAME OSTEOPENIA) (TYPE BEGIN-INTERVAL)
      (REFSET (OSTEOPOROSIS)))
      ((REFSYSFORM 'MENOPAUSE-ONSET') (REFSYS DEVELOPMENT)
      (TYPE POINT))
      (-5 YEARS) (0 YEARS))
```

Definitions:

- *Osteopenia* → decreased bone density.
- *Rheumatic fever* → an acute disease that follows infection with specific types of the streptococcal bacterium sometimes leading to cardiac involvement.
- *Rheumatic heart disease* → heart disease caused by rheumatic fever.
- *Dyspnea* → shortness of breath. One of the many etiologies for dyspnea is heart disease.

Given two RRELs r_1 and r_2 such that $i \xrightarrow{r_1} j$ and $j \xrightarrow{r_2} k$, reference sets modify constraint propagation by only calculating the range addition of $i \xrightarrow{r_1} j$ and $j \xrightarrow{r_2} k$ if the intersection of the reference set membership of points i, k is non-null. This

⁷These temporal assertions are meant to describe events in an individual patient model, and not a population hypothesis.

effectively limits constraint propagation to single reference sets unless there is a significant degree of overlap between reference sets. Even then, as shown in Figure 2.14, propagation can only spread from one reference set to another if an RREL within the overlapping area is constrained. Note that points within reference sets are exhaustively interconnected by RRELS.

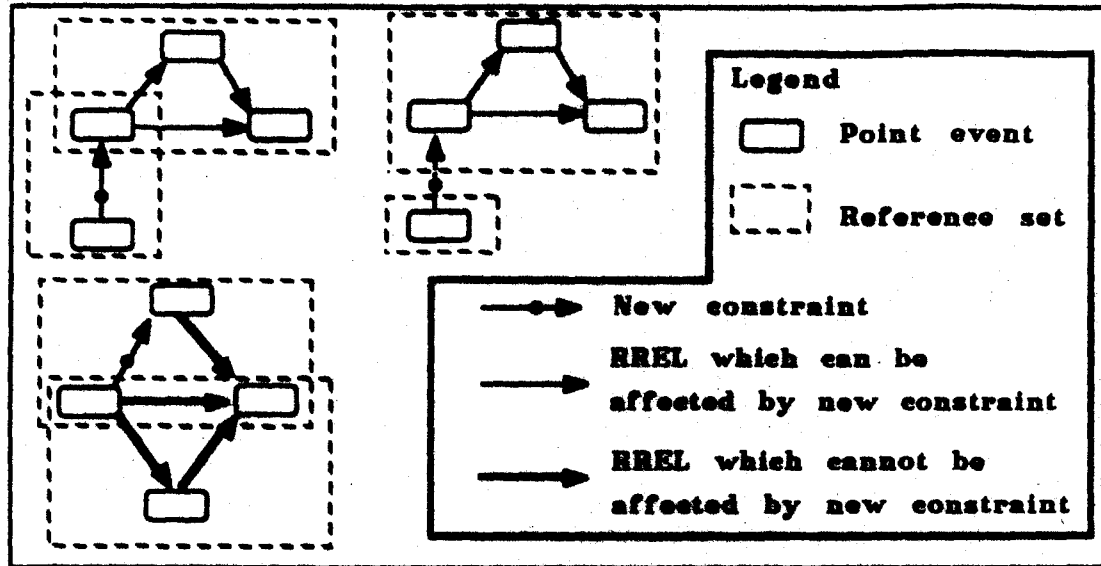


Figure 2.14: Effect of Reference Set Configuration upon Constraint Propagation.

Once the temporal data base is clustered into reference sets, the relation between two points can no longer be simply obtained by retrieving the one RREL that directly links the two. Instead TUP must use a variant of the "best-first" graph search. In TUP, the best-first search attempts to find that path between the two points specified that provides the most constrained values for the temporal relation. This search is guided by a cost evaluator that determines which node expansion (i.e. which RREL to select) will give the minimum cumulative cost.

The A^* search [19] could be implemented if, in addition to cost of the path searched so far, the cost evaluation function would estimate the cost of the path remaining to be traveled. In the current implementation, only the first part is calculated (by repeated range-additions). A first cut at obtaining a function that estimates the cost of the path remaining to be traversed would be to make very large cost estimates for those paths that traverse reference sets that do not connect

to the goal point.⁸ Another, unimplemented, method for increasing the efficiency of the search would be to make it bidirectional—starting searches from both points and meeting in between [20].

When the search explores a reference set, it only requires one node expansion because of the exhaustive interconnections within reference sets. Also, by definition, it never takes more than the traversal of one RREL to move from one reference set to another. Therefore, in the worst case, the maximum depth of the search is twice the number of reference sets.

2.4.2 Costs of Clustering

Clustering the temporal data base does incur some costs in terms of the quality of temporal information retrieved. These costs are minimized by judicious selection of the reference sets. Discussion of the criteria for the selection of reference set membership is deferred to section 2.4.4.

As RRELs in separate reference sets do not affect⁹ their respective state of constraint, an RREL wholly within a reference set is no longer guaranteed to be maximally constrained. In practice, if the criteria for the selection of reference set membership are followed closely, assertions of RRELs within a reference set will contribute most, if not all the constraints within the reference set, and consequently little is lost by ignoring temporal relations outside the reference set. Even so, if the RREL retrieved is deemed insufficiently constrained by the user or expert system, a best-first search can be initiated¹⁰ that will use information throughout the temporal data base, rather than just within the reference set.

For the same reasons that maximum constraint can no longer be guaranteed for RRELs contained within a single reference set, no guarantee can be made for data base-wide temporal consistency. However, since constraint propagation within reference sets does ensure local consistency, many inconsistencies will be detected.

In the unlikely event that an undetected inter-reference set inconsistency occurs, there is no easy solution short of eliminating the reference sets and permitting unrestricted constraint propagation. Even in this rare eventuality, I think this inconsistency will not be of major consequence to the performance of the host

⁸ A precompiled connectivity map of reference sets could be used for this purpose.

⁹ Except occasionally in the cases of overlapping reference sets discussed above.

¹⁰ Recall that the search is automatically performed if the temporal relation desired is of two points with non-intersecting reference set membership.

expert system. Using the most recent example, if one of the RRELS within the rheumatic heart disease reference set were to be inconsistent with the RRELS in the osteoporosis reference set, this probably would not impede diagnosis of either of the two conditions as long as the reference sets were internally consistent (which is guaranteed by constraint propagation).

2.4.3 Computational Savings

What then are the computational savings that one can expect from the use of reference sets? With two hundred point events in a hypothesis, unrestricted constraint propagation would create approximately twenty thousand RRELS. The same two hundred points clustered into twenty reference sets of ten points each would generate approximately one thousand RRELS. That is, the number of RRELS grows linearly with the number of events rather than with the square of the number of events.

The space savings are much less than the savings in the time required for propagating all the constraints since the computation time for constraint propagation increases with the cube of the number of events. With reference sets of uniform size the time for constraint propagation grows linearly with the number of events.¹¹

The impact of these savings can be appreciated by considering the performance of TUP with a hundred point events. Without reference sets, the space capacity of the machine¹² was exceeded after more than one hour of computation whereas with ten reference sets, of ten points each, the time for constraint propagation was 45 seconds.

2.4.4 Design of a Temporal Clustering Heuristic

The overriding motivation in the design of a clustering heuristic was that it reflect the goal that led to the creation of clusters in the first place: computational tractability. In the context of constraint-propagation-based reasoners this means that relations that are retrieved often should be precomputed whereas infrequently retrieved relations can span clusters and therefore require search. This clustering by frequency of retrieval differs markedly from schemes that cluster by temporal location. Although temporal proximity or contiguity may generate clusters with

¹¹Of course, these savings will be nullified if many of the queries, subsequently put to TUP, require inter-reference set search.

¹²A XEROX 1108 running INTERLISP D.

the desired properties, there are many domain applications where this will not be the case.

With most clustering schemes, if the knowledge engineer is not careful, too many of the relations retrieved will require the relatively expensive search procedure. In the case of Allen's reasoner, since relations between intervals within the same reference interval are computed locally, there may exist inference paths that include interval relations outside the reference interval that would infer a more constrained relation. Similarly, inconsistencies that would be detected by the original constraint propagation algorithm might be overlooked. The knowledge engineer must therefore be careful not to omit from a reference interval those intervals and relations that might increase the constraint of the relations contained. Deciding *a priori* which of the relations will be retrieved frequently is, as Allen suggests, a difficult task. It requires that the knowledge engineer understand and anticipate much of the dynamics of the performance program. In a medical expert system this involves knowing which temporal relations might be pivotal in differentiating hypotheses.

Take for example the case of a sixty-year-old woman who has received a blood transfusion and has jaundice. An incidental finding of osteopenia is noted on the basis of an abdominal X-ray taken during the work-up. A medical expert or expert system should consider the possibility of a blood-borne viral liver infection and in differentiating this hypothesis from other causes of jaundice, the temporal relation between the transfusion and the onset of jaundice will be much more important than the relation between the jaundice and osteopenia. The former will be therefore retrieved more frequently for comparison with other hypotheses than the latter. Fortunately, in many applications much information about the mutual importance of events is already encoded in the knowledge structures. This knowledge engineering practice was advocated by Bobrow and Winograd [6, page 267]:

One of the fundamental problems in artificial intelligence is the "combinatorial explosion." A large knowledge base provides an exponentially expanding set of possible reasoning chains for finding desired information. We believe that the solution to this problem must be found by dealing with it directly through explicit concern with the *accessibility* of information. The representation language must provide the user with a set of facilities for controlling the way in which memory structures are stored, so that there will be a correspondence between "salience" or "relevance" and the information accessed by procedures for search and reasoning operating under processing resource limitations.

Even if such correspondence is not always a conscious concern of the designers of knowledge representations, the demands of computational tractability will frequently push development in that direction. Otherwise, the combinatorial explosion, mentioned above, will grossly impair the performance of the applications that use the representation.

The parallel between salience and the accessibility of information in knowledge structures can be exploited by temporal reasoners to create clusters or cluster hierarchies that parallel those of the atemporal structures. The Salience Clustering Heuristic (SCH) does just that.

To obtain the expected improvement in performance with SCH, the knowledge structures must fulfill the criteria listed below. These are illustrated in the example in Section 2.4.1.

- **Parallel salience.** There must be correspondence between atemporal and temporal salience. That is, in the particular application being considered, there must be some explicitly represented atemporal organizing principle that groups together the events linked by important (frequently retrieved) temporal relations. Although in most planner, object-oriented, and process representations such correspondence exists, this is not necessarily the case.
- **Cluster Size.** The maximum cluster size depends on the particular flavor and implementation of temporal inference, the machine it is implemented on, and the maximum time for temporal inference that is acceptable to the user of the performance program.
- **Disjointness.** If there is too much overlap between clusters, and constraint propagation frequently spreads through the areas of overlap, then the effective cluster size will be significantly increased. SCH therefore works best when the clusters are mostly disjoint.

2.5 Automatic Generation of Reference Sets

THRIPHT uses SCH to derive reference set membership from the causal aggregation hierarchy used in representing each hypothesis (Figure 2.15, page 47). The implementation details are deferred to chapter 3, page 82. Here, I discuss whether the causal aggregation hierarchy is an adequate substrate for SCH by examining the applicability criteria itemized above, but first, a few definitions are in order.

Causal links, as used here, are skeletal versions of those used in the ABEL program. That is, a causal link is a relation between two events such that the onset of the causally antecedent event is simultaneous with, or before the onset of the causal consequent and the likelihood of the causal consequent occurring is increased if the causal antecedent occurs.¹³ Although a truly usable implementation of a diagnostic medical expert system would require a more powerful representation of causal association, I have found this definition sufficient to demonstrate the role of temporal reasoning in causal-association-based systems.

A causal aggregate is a summary of the description of a disease course at a more detailed level. As in ABEL, the more detailed level is represented by a network of events, which themselves may be causal aggregates, interconnected by causal links.

Parallel salience: Each causal link, by its nature, has an associated set of temporal constraints, minimally that the onset of the effect follows, or is simultaneous with the onset of the cause. For reasons of computational parsimony and epistemic sufficiency, only the relevant causal links are explicitly represented. In this application, medical diagnosis, the important causal links are those that are characteristic of disease hypotheses. The time constraints between cause and effect are intrinsic to this characterization. Since the organising principle—causal aggregation—encapsulates networks of causal links within an abstracted description these aggregated links must have some mutual relevance. The parallel between the relevance of causal links and temporal relations ensures that temporal clusters based upon causal aggregation will group mutually relevant temporal relations.

An example: as diagrammed in Figure 2.15, Inoculation and Jaundice will fall within the reference set *Prodrome/Acute-Hepatitis*. The temporal relationship between the two is important for arriving at the diagnosis of hepatitis B since it permits a whole host of other hypotheses to be dismissed (or at least ranked low in the differential diagnosis). In the case of a patient with a history of transfusion (the putative inoculation) followed by prodromal signs and symptoms (e.g., malaise, chills and fever), if the delay between transfusion and jaundice is between 50 days and 7 months (a range corresponding to the viral incubation period) the hypothesis of blood-borne hepatitis B gains weight. If the prodromal symptoms follow the transfusion by couple of days, a blood-type mismatch is much more likely.

In contrast, the temporal relation between the inoculation and the onset of recovery (which do not share reference sets) is considerably less useful for diagnostic

¹³i.e. $p[\text{consequent}] < p[\text{consequent}|\text{antecedent}]$




Figure 2.11: Triggered hypothesis for hepatitis B

purposes because it is so variable (and can extend to a lifespan if the patient develops chronic active hepatitis). This illustrates another property of parallel salience: that temporal relations within clusters are usually more constrained than those between clusters. Why this is so is unclear: perhaps more effort is invested in obtaining and asserting more precise estimates of the temporal relationships that are characteristic of a disease and on which diagnosis might hinge. Alternatively, events that are closely related to a particular causal mechanism might have less variable timing with respect to one other than with respect to events that are only indirectly related to that mechanism.¹⁴ In any case, this is useful as it means that most of the constraint in the computation of an REEL can be obtained locally within a cluster. This helps minimize the loss of information that Allen mentions regarding reference intervals.

Cluster Size: The number of events that are included within a reference set depends upon the number of levels of the causal aggregation hierarchy that SCH uses. With unabbreviated medical hypotheses, each causal aggregate has five to fifteen elements. In the current implementation of TUP this is at the upper bound to achieve anything approaching real-time performance. Consequently, THIRPHT directs SCH to gather only the immediate elaboration of a causal aggregate within a reference set.

Disjointness: Within a particular hypothesis, only one or two events are shared by causal aggregations. The effective cluster size is therefore approximately the same size as the reference set membership.

It seems then that events with the same *immediate* causal aggregation share the criterion of relevance necessary for membership within the same reference set. Note that temporally distant events (as in our earlier example of rheumatic fever and mitral stenosis) may share the same immediate causal aggregate and therefore share reference sets. Also, the component events of a causal aggregate are not necessarily DURING the aggregate. For instance, the immune response to a Hepatitis B infection persists into the recovery period.

2.5.1 Using Other Knowledge Representations

Most if not all knowledge structures are built with some regard to processing resource limitations and thus indirectly to the salience of the entities within the knowledge structures. This includes the whole spectrum of knowledge representations: causal aggregation hierarchies [45,50], structured object representations [6,48,40]

¹⁴i.e. In which there are several intermediary mechanisms.

goal structures of planners [58], qualitative system descriptions [17,29,13] and the more recent hybrid knowledge representations such as KRYPTON [7] and KL-TWO [60]. When, as in the previous example from THIRPHT, temporal salience follows atemporal relevance, SCH can be brought to bear. I outline below some of the promising candidates to which SCH can be applied.

2.5.2 Qualitative Reasoning

In qualitative simulations [17,30,13] the number of events, and associated temporal relations, grows rapidly with the number of processes modeled, especially if there is significant interaction between these processes. Williams's [64] approach to making temporal queries computationally tractable is to generate *justification histories*. These histories maintain the dependencies between the quantity values at different times. Temporal reasoning can therefore be safely restricted to those events that share membership in a justification history. Nevertheless, in large simulations such as models of human physiology, the extent of interaction between processes will create justification histories that again contain too many events for reasonable performance in temporal reasoning.

SCH, however, can exploit the effort already invested to "chunk" process descriptions *a la* Qualitative Process Theory [17] by using the *parameter histories* [17, page 28] associated with the Individuals in a process description (Figure 2.16) to determine reference set membership. By definition, the Individuals (and associated parameters) in a process description (*s,d* and path in the example) have greater shared relevance¹⁵ than they have with other individuals in other process descriptions.

What about SCH's requirement for parallel temporal and atemporal salience? Clearly, this depends on the application so let us take a typical medical application of qualitative simulation: explanation of the causal mechanism of pathophysiological phenomena [29].

An adequate explanation of the causal behavior of a system must include a partial order, if not quantitative chronology, of the events in the behavior. It is especially important, if the explanation is to be consistent, that the timing of events which are tightly coupled (closely related to a shared mechanism) be as precise and consistent as possible. It is less important to ensure precision and consistency

¹⁵In describing a particular mechanism.

```

process: fluid-flow
Individuals:
  s a contained-liquid
  d a contained-liquid
  path a fluid-path, Fluid-Connection(s,d,path)
Preconditions:
  aligned(path)
QuantityConditions:
  A[Pressure(s)] > A[Pressure(d)]
Relations:
  Let flow-rate be a quantity.
  flow-rate PROPORTIONAL-TO A[Pressure(s)] - A[Pressure(d)]
Influences:
  I+(Amount-of(d), A[flow-rate])
  I-(Amount-of(s), A[flow-rate])
; A fluid path is aligned only if either it has no valves
; or every valve is open

```

Figure 2.16: Process Description of Fluid Flow.

between events that are less tightly coupled or related to completely separate mechanisms. For instance, there is a greater need for precision in describing the timing of decreased oncotic pressure¹⁶ and subsequent increase in transudation¹⁷ than for the timing of decreased oncotic pressure and the onset of a low protein diet.¹⁸ Precision in the timing of decreased oncotic pressure and the tremor events of Parkinson's disease is even less important as their respective mechanisms are distantly, if at all, related. In this last case, timing information would not contribute to the understanding of the dynamics of either mechanism.

Observe that we seem to have more precise temporal information (RREs with a lower constraint index) about tightly coupled events than about loosely coupled or unrelated events. As in the case of causal-association networks, this is likely due to

¹⁶This usually refers to the osmotic gradient, across a semi-permeable membrane, set up by constituents of a medium such as blood or interstitial fluid.

¹⁷The flow of an ultra-filtered fluid across a semi-permeable membrane consequent to an osmotic or hydrostatic gradient.

¹⁸Which, in severe cases, can lead, through several intermediate mechanisms, to decreased oncotic pressure.

the greater effort invested in obtaining and asserting more precise estimates of the temporal relationships of those events which are tightly coupled to, and therefore characteristically associated with the behavior of a mechanism. Or, each additional intermediate mechanism, between two events, adds an increment to the variability of the temporal relation between those events.

The parallel between temporal and atemporal salience permits SOH to take the union of the events of the parameter *histories* associated with the Individuals in a process description and lump them in a reference set so that their temporal relations are precomputed by constraint propagation while relations between events in different processes are obtained by search. Those processes that interact, and therefore share parameter histories, have correspondingly overlapping reference sets. As in most applications of SOH, the usual caveats regarding potential loss of information apply.

2.5.3 Planners

Planners have the properties of parallel salience and disjointness that make clustering effective. That is, most of the temporal relations asserted within a plan are between actions sharing the same proximate supergoal—a major reason Vere's clustering scheme works. Application of SOH to a planner goal hierarchy simply involves declaring a goal to be a reference set of the subtree of goals it subsumes. The depth of the subtree is determined by the cluster size criterion. That is, SOH is applied recursively until the number of planned events in each subtree is within an acceptable range.

2.5.4 Frame-based and Hybrid Representation Languages

Languages such as KRL [6] and KI-ONE [8] have the breadth of expression to provide many taxonomic distinctions with which to drive SOH. Whether this taxonomical structure can be used in a particular application depends on the extent to which the criteria for SOH are met. This must be determined by the knowledge engineer.

Take KRL's perspectives: from one perspective an individual can be viewed as a traveler and from another as a customer. By definition, the components of a perspective share a common view or relevance. Temporal salience may follow that imposed by a perspective, for example: the timing of events in the traveler perspective (e.g. geographical movements) have greater mutual salience than they have

with customer perspective events (e.g. book purchases). If the reference sets generated by perspectives are too large they can be subdivided using other taxonomic structures of the language such as the SetOf set membership specification.

KL-ONE and its derivatives attempt to distinguish between the descriptive, terminological component of a representation and the assertional component. KRYPTON implements the former as the TBox and the latter as the ABox. The TBox is used to build descriptions of the world, "the formal equivalent of *noun phrases* such as 'a person with at least three children' " [7, page 418] from which the logical consequences (such as subsumption and disjointness) of assertions made in the ABox can be retrieved. The TBox provides a rich terminology with which to describe the roles that concepts fill and the manner in which concepts are specialized. This terminology provides many dimensions with which to group concepts with shared salience. For instance, concepts which fill the same role

Unlike process descriptions and causal-link-based hypotheses, these last representations are much more general and domain independent. The parallel between the atemporal clustering and the shared salience of the temporal relations is domain dependent and may not always hold. Therefore, a knowledge engineer cannot borrow knowledge structures from other applications and hope to effectively use SCH without first verifying that the parallel holds in the new application.

Performance Driven Clustering

All the above cases use domain-dependent knowledge of the salient decomposition of the domain knowledge to drive the clustering of the temporal data-base. Could this clustering be done automatically and yet wholly internally to TUP without knowledge of the domain?

At the beginning of the discussion on clustering the temporal data-base, I asserted that creating a cluster on the basis of a purely temporal criterion—especially temporal proximity—was unhelpful as it did not capture any notion of mutual relevance of the member events of the cluster. However, in the course of my work in this area, it has become apparent that TUP has access to several clues to the domain-dependent salient decomposition even while restricting itself to introspection upon its own operations. First, the frequency of retrieval and assertion of certain RRELS

relative to others in the data-base provides strong guidelines for the correct¹⁹ reference set configuration. Second, the pattern of distribution of the constraint index would also point to the correct reference set decomposition; those connected²⁰ portions of the data-base possessing a significantly lower CI should be assigned to the same reference set.

A version of TUP, modified to perform temporal clustering in the manner described above (performance-driven clustering), would start with a collection of assertions upon which no constraint propagation had been performed. All retrievals at this initial stage would use the search mechanism. After additional RRELS would be asserted, performance would degrade as the search paths became progressively longer. Upon reaching a threshold of poor performance, the performance-driven-clustering version of TUP would bring to bear several heuristics for clustering the data-base into reference sets. Once the reference sets were defined, constraint propagation would be initiated within each reference set. The boundaries of these reference sets would be re-evaluated by the heuristics as additional data were accumulated and more queries made. As the configuration of these reference sets approached the natural decomposition of the knowledge in the application domain, the reference set configuration would become increasingly stable.

The heuristics that would be employed would have to be precisely tuned so that the criteria for reference set membership were neither too stringent nor too lax. The high-level strategies would, however, be:

- If a group of connected, externally asserted RRELS have an individual assertion rate x times higher than other RRELS to which they are connected, then make a reference set out of the members of this group.
- If a group of connected, externally asserted RRELS have an individual retrieval rate y times higher than other RRELS to which they are connected, then make a reference set out of the members of this group.
- If a group of connected, externally asserted RRELS have an individual constraint index z times less than other RRELS to which they are connected, then make a reference set out of the members of this group.

¹⁹Recall that RRELS within reference sets are much cheaper to retrieve than inter-reference set RRELS and therefore it is desirable to have frequently asserted and retrieved RRELS within a reference set.

²⁰Points connected by externally asserted RRELS.

- If a proportion n of the reference sets have memberships greater than m events, modify the thresholds of the first three heuristics to make reference set membership more stringent.
- If a proportion i of the reference sets have memberships less than j events, modify the thresholds of the first three heuristics to make reference set membership more lax.

I find this scheme of performance-driven-clustering quite intriguing, because it creates a temporally-oriented memory that progressively organizes itself so that those temporal relations that connect events that are closely related are those that are the most easily retrieved. It also guarantees temporal consistency for those temporal relations that connect closely related events. In contrast, temporal relations that are rarely retrieved, between events that are not very relevant to one another, require considerable effort (search) for retrieval.

2.6 Reference Systems and Reference Sets

The earlier description of the commonly used temporal yardsticks (*reference systems*) such as the calendar, was purposefully left incomplete pending the discussion of constraint propagation and the role of reference sets. I omitted the fact that when the mini-expert of each reference system asserts temporal relations between members of a reference system, constraint propagation causes the exhaustive interlinking of all members of this reference system. This constraint propagation is restricted to the reference system by the creation of a reference set that contains only members of the reference system. This reference set is created by the mini-expert which appends to each point asserted a reference set whose name is that of the reference system as in example 12.

Example 12

```
(RREL ((REFSYS CALENDAR)
      (REFSYSFORM "8 p.m., May 20th, 1986")
      (REFSET CALENDAR))
      ((REFSYS CALENDAR)
      (REFSYSFORM "10 p.m., May 20th, 1986")
      (REFSET CALENDAR))
      (2 HOURS) (2 HOURS))
```


These mini-experts have knowledge of the temporal distance between points in a particular reference system and are implemented as black-boxes that only provide TUP with the distance between two points, specified in canonical form. The internal operation of these black-boxes is quite varied, the CALENDAR mini-expert uses a simple formula to obtain temporal distance between points whereas the DEVELOPMENTAL mini-expert has its private temporal context with range relations from which it obtains this distance information.

Early in TUP's development, one of the problems that arose was the inability to use reference system information from within (non-reference system) reference sets without having recourse to a search. In Figure 2.17 (a) the RREL linking points a and b is less constrained than the RREL that could be obtained if a search which included the reference system information (along the path $a \rightarrow d1 \rightarrow d2 \rightarrow b$) was performed. Nevertheless, since retrieval of an RREL between two point events in the same reference set is done directly and without search, the reference system information is not employed.

In general, one would like to be able to always use reference system information to constrain RRELS within reference sets because these reference systems are ubiquitous, and particularly because the constraint index of these reference systems is usually very low. The CALENDAR reference system, for instance, has a total constraint index of zero as all relations within that system are maximally constrained.

The solution adopted involves adding points of a reference system to the reference set of those events to which they are related by external assertions (Figure 2.17 (b)). By dint of constraint propagation, the RREL between a and b is then appropriately constrained by the calendar information. Note that because of the restrictions on constraint propagation imposed by reference sets, range addition is not attempted between RRELS exclusively in the X reference set and those exclusively in the CALENDAR's reference set. Therefore the only way that constraint could propagate from the X reference set to the CALENDAR's reference set would be if an RREL within the region of reference set intersection were to be constrained.

As reference systems have a low CI, the direction of constraint propagation is usually from the reference system to the local reference set. Only rarely does it occur that RRELS within a reference set constrain relations in a reference system. When it does, as might happen for instance, if a reference set were to contain an assertion of early DEATH then one would want all reference sets that contained events bounded by DEATH to be appropriately constrained. This in fact is what would happen under the rules of constraint propagation as modified for reference sets.

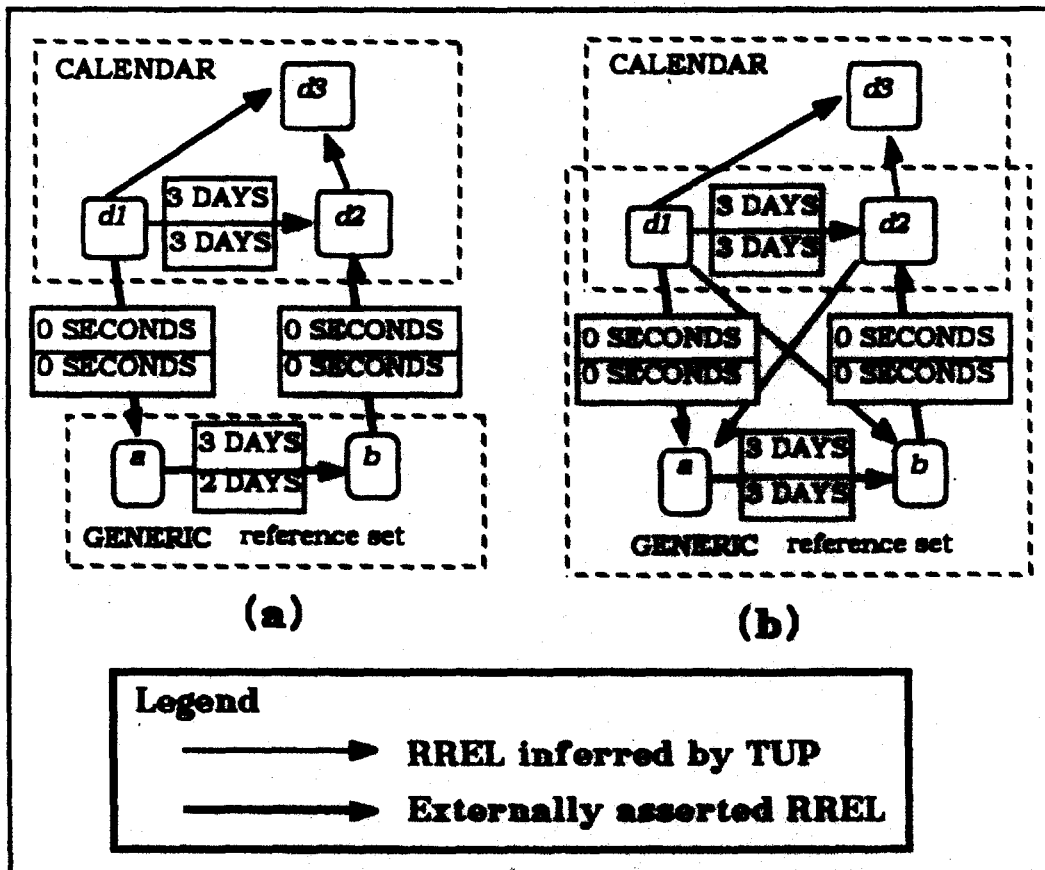


Figure 2.17: Use of Reference System Information.

Note that the solution adopted remains true to the *leitmotif* of reference sets—that only closely related event points be included in the same reference sets. In this case, only those points in the reference system that are externally asserted to be related to points in a reference set are included in that reference set.

2.6.1 The PRESENT Reference Set

Assertions of the relationships of events with respect to the present are typically found throughout the temporal data-base, involving events of several different reference sets. Also, as most instances of the present are generated by means of the RelationToPresent assertion, most NOW events will be directly linked to a time and date. Consequently, just as reference systems are a source of tight constraint, so are references to the present. Therefore, the same solution adopted for reference sets was implemented for NOW instances. That is, all instances of NOW are made (by RelationToPresent) to be members of the PRESENT reference set. Also, as for reference systems, the same instance of NOW is made a member of the reference set of the point specified in the RelationToPresent assertion. Also, by the mechanism described in the previous section, the reference set information of this NOW point is appended onto the reference set membership of the CALENDAR point that represents the current time and date. The final result is illustrated in example 13, the complete version of example 6.

Example 13

```
(RelationToPresent
  ((NAME ASTHMA) (TYPE BEGIN-INTERVAL)
   (REFSET (RESPIRATORY-DISTRESS)))
 (+12 MONTHS) (+13 MONTHS))
```

translates to:

```
(RREL ((NAME (GENSYM NOW)) (TYPE POINT)
       (REFSET (PRESENT RESPIRATORY-DISTRESS)))
      ((NAME ASTHMA) (TYPE BEGIN-INTERVAL)
       (REFSET (RESPIRATORY-DISTRESS)))
      (-13 MONTHS) (-12 MONTHS))

(RREL ((REFSYS CALENDAR)
```

```
(REFSYSFORM (DATE)) (TYPE POINT)
(REFSET (CALENDAR PRESENT RESPIRATORY-DISTRESS)))
((NAME (LatestNow)) (TYPE BEGIN-INTERVAL)
(REFSET (PRESENT RESPIRATORY-DISTRESS)))
(O SECONDS) (O SECONDS))
```

2.7 Predicates and Retrieval Functions

Up to this point, I have described much of TUP's functionality, without touching upon the practical aspects of using TUP in a medical expert system. Although the RREL can be used to assert and retrieve all temporal information, the direct use of the RREL form in most applications is at the best cumbersome. For this reason, TUP's implementation includes a broad variety of predicates for testing different kinds of temporal relationships and a few specialized retrieval functions.

2.7.1 Predicates

The nature of the RREL makes for two kinds of tests of temporal relations. One kind determines whether the specified conditions could possibly be true while the other kind tests whether these conditions are *strictly* true. For each TUP predicate, there consequently is a corresponding strict²¹ and a relaxed version. Take the WITHIN-P predicate, it tests whether two time points fall within a specified distance of one another as in form 7. In the strict version (S-WITHIN-P), the predicate returns "true" only if both the upper and lower bounds of the RREL between the two time points are less than or equal to the specified distance. The relaxed version (WITHIN-P), will return "true" as long as one of the two bounds is less than or equal to the specified distance. Subsequent constraint propagation may further constrain the RREL such that the WITHIN-P predicate might return "false." By definition, as S-WITHIN-P is the strict version of the predicate, it never changes the value it returns from "true" to "false", no matter how much the tested RREL is constrained. In practice, the relaxed predicate versions are the pragmatic choice for knowledge engineering in application domains where temporal knowledge is poorly constrained. All the implemented TUP predicates are described in appendix A.

Form 7

```
(WITHIN-P <point 1> <point 2> <distance> <context>)
```

²¹Denoted by an "S" prefix to the predicate's name.

2.7.2 Event Retrieval

TUP implements event retrieval (as opposed to RREL retrieval) with the GETEVENT function which has the form:

Form 8

(GETEVENT <point specification> <filter> <context>)

where the filter is a boolean combination of TUP predicates. If the point specifications (e.g. reference set membership) match a point (or points), and the filter returns "true", then the point(s) matched are returned by the function. This permits the retrieval not only of events of a particular type or name, but also with a specific temporal relationship with respect to other events.

FINDBETWEEN and FINDPOSITION

There are some applications in which one is interested in the event(s) that occur during an interval²² and the ordinal position of an event in a set of events. The FINDBETWEEN and FINDPOSITION functions (see form 9) provide this capability.

Form 9

(FINDBETWEEN <point 1> <point 2> <scope> <context>)

(FINDPOSITION <point> <point set> <context>)

The FINDBETWEEN function obtains the intersection of all the events returned by GETEVENT that are after the first point, with all the events before the second point. There is a strict and relaxed version of this function, the only difference being whether strict or relaxed BEFORE-P and AFTER-P predicates are employed; the relaxed version of FINDBETWEEN tends, of course, to return a greater number of events.

²²Be it either a specific interval on the calendar reference system or simply the interval between two point events.

When someone asks, "What has happened today?" your answer will depend on what category of events you believe the questioner is referring to. To a doctor you might enumerate the various symptoms you have experienced, while to a business associate you might enumerate the transactions of the day. The scope parameter of the `FINDBETWEEN` function permits the specification of a list of reference sets to which the function should be applied. In addition to delimiting the categories that `FINDBETWEEN` will be applied to, this scoping reduces the expense of a very expensive computation. Of course, if necessary the scope can be specified to include the whole temporal knowledge-base.

`FINDPOSITION` returns the ordinal position of a time point in a set of time points. It does this simply by counting the number of events in the set that are `BEFORE-P`²³ the specified point. There is no need for scoping as the scope implicitly only includes the points that are members of the point set. I have found however, that in most realistic, large applications, the point set is too large to be directly specified, but instead is used with the `FINDBETWEEN` function as in example 14.²⁴ Therefore, in practice, the `FINDPOSITION` function incurs approximately 50% more computational expense than `FINDBETWEEN`. As discussed in chapter 5, this phenomenon has its parallel in human cognition.

Example 14

```
(FINDPOSITION
  (GETEVENT ((NAME FEVER) (TYPE BEGIN-INTERVAL))
    (S-AFTER-BY-P
      ((NAME FEVER) (TYPE BEGIN-INTERVAL))
      ((NAME INOCULATION) (TYPE BEGIN-INTERVAL))
      (4 DAYS) (10 DAYS)))
  (FINDBETWEEN ((REFSYSTEM "7-4-84") (REFSYS CALENDAR))
    ((REFSYSTEM "7-4-85") (REFSYS CALENDAR))
    *ALL-DISEASE-REFERENCE-SETS*))
```

Although these functions are indeed very expensive relative to other TUP retrieval operations, it is my experience, at least in the domain of medical diagnosis, that they are infrequently needed.

²³The strict version `S-FINDPOSITION` that uses `S-BEFORE-P` is also implemented.

²⁴Note the use of the filter by the `GETEVENT` function as well as the scoping of the `FINDBETWEEN` function. Context slots have been left at their default value.

2.8 Persistence

Any event in THRIPT has a persistence that is bounded by the interval in TUP to which it is associated. Often there is some initial knowledge of the extent of this persistence (e.g. we know *a priori* that cardiac angina without ischemia does not last several hours) that is an intrinsic property of the event. Also, the persistence of an event is bound by the persistence of other events (e.g. the duration of diabetes mellitus is limited to an individual's lifespan). This is represented in TUP by having the end of one interval tied (with an RREL) to another. The mutual restriction of persistence that this interdependence represents is taken care of by constraint propagation. For instance, every time the lifespan duration is shortened, so is the persistence of diabetes mellitus.

There are, however, many events whose persistence remains completely unknown. TUP makes the same assumption in this regard that a lot of people would make: that the persistence extends to infinity—at least until further knowledge is gained about the intrinsic persistence of such an event or the persistence of another event to which it is tied. For example, the duration of the interval that represents the persistence of the planet Earth's existence will have an upper bound of $+\infty$. To hedge the bet, there could be another event—the destruction of the Earth—whose corresponding interval would have an onset simultaneous to the end of the existence of the planet. The moment the discovery was made that planets such as the earth had a specified limited lifespan or that there would be planetary destruction in less than infinite time, the infinite persistence of Earth's existence would be constrained accordingly.

2.9 Unresolved Issues

TUP has been designed to express temporal knowledge as generally as possible while maintaining a simple underlying temporal representation. Even so, TUP's performance falls short in the representation of certain kinds of temporal information, specifically: recurrent events, probabilistic distribution of temporal bounds and parametrisation of bounds.

2.9.1 Reasoning With Temporal Uncertainty

Several temporal reasoning systems, capable of hypothetical reasoning through the use of contexts or some form of backtracking (or both) are billed as endowed with the ability to reason with temporal uncertainty. In some sense this claim is correct, but only in a restricted sense—as will be explained shortly.

Let us see what kind of reasoning about temporal uncertainty is permitted by context and backtracking mechanisms. If the temporal representation is used in an *assertional* mode, a backtracking mechanism permits the temporal reasoner to correct prior assumptions as additional, possibly unexpected, information is obtained. In this sense, the temporal reasoner is handling uncertainty as it responds to events whose occurrence cannot be fully anticipated. If the temporal representation is used in a descriptive or *terminological* manner, alternate, distinct temporal hypotheses can be described as well as the logical (temporal) conclusions that derive from specific temporal relationships. Here again, the temporal reasoner can be described as reasoning about temporal uncertainty, because it handles the uncertainty over the “real” timing of events by representing several different possible temporal configurations.

THRIPHT exhibits both the assertional and terminological methods of dealing with uncertainty in its use of TUP's reasoning mechanism and representation. In the data-gathering (first) phase of history-driven-diagnosis, inconsistencies detected by TUP cause the withdrawal of one or more RRELS and the subsequent recomputation of the consequent temporal relations. The terminological representation of uncertainty is illustrated by THRIPHT's alternate causal/temporal hypotheses generated in the elaboration phase.

Encoding Probabilistic Information

Both of the above methodologies for representing temporal uncertainty are too weak to express more detailed temporal knowledge. In particular they do not allow an expert system to make use of probability estimates of temporal relations. In this section, I sketch some of the challenges of providing such a capability.

Let us take two RRELS used in previous examples, and suppose that we could assign a prior probability of p_1 to the temporal relationship of the first RREL and p_2 to the second RREL (see example 15). What would be the probability of the third RREL linking onset of anorexia and end of irritability? To begin with, the question itself is flawed, because the semantics of RRELS are such that the RREL bounds between

two point events are not determined probabilistically, but are the fixed limits for that temporal relation in a specific TUP context. Any other interpretation renders the results of the range addition operation meaningless.

Example 15

```
(RREL ((NAME IRRITABILITY) (TYPE BEGIN-INTERVAL))
      ((NAME ANOREXIA) (TYPE BEGIN-INTERVAL))
      (2 DAYS) (3 DAYS) (Prior: P1))

(RREL ((NAME ANOREXIA) (TYPE BEGIN-INTERVAL))
      ((NAME ANOREXIA) (TYPE END-INTERVAL))
      (72 HOURS) (96 HOURS) (Prior: P2))
```

We can, however, redefine the semantics of the RREL so that in our example, the interpretation of the first RREL becomes: "There is a probability p_1 that the temporal distance between the onset of irritability and the onset of anorexia lies in the range of two to three days." The operation equivalent to range addition should then either calculate the bounds of the third RREL at a pre-set probability p_3 or calculate the probability p_3 for a particular pair of bounds. Either of these two operations really requires not just a single probability estimate, but in fact the probability distribution of the temporal distance for each pair of event points. Since the probability distribution of a calculated RREL depends on the probability distributions of the two RRELS used for the calculation, this capability would require that TUP be capable of performing the equivalent of range addition on any pair of arbitrary probability distributions.

Let us go a step further, and imagine that the probability distributions, over the full range of temporal distances, were represented for each RREL. Just as in the case of the range addition calculation, there would be several combinations of RRELS that could be used to calculate a particular probability distribution. The problem would then arise of how to resolve the differences between the probability distributions calculated from different RREL combinations.

Although this is a difficult problem, in the absence of a solution, the manner in which temporal reasoners deal with temporal uncertainty will remain, at best, *ad hoc* and subject to inconsistent interpretations. In this respect, the recent work in mathematical reasoning with partially specified systems of equations²⁶ appears to be a promising line of research. If we could provide the reasoner with the general

²⁶As exemplified by Sacks' [53] Qualitative Mathematical Reasoner.

form of the distribution (e.g. normal distribution) as well as a few values along this curve (e.g. 75% of individuals have symptom x within five days of symptom y), this might sufficiently constrain the set of consistent probability distributions to the point that useful temporal calculations could be made with them. Note, however, that in medicine at least, temporal information of quality sufficient to construct a temporal distribution is scarce,²⁶ and therefore probabilistic temporal reasoning would be infrequently useful.

2.9.2 Recurring Events

Due to the significant difficulties in providing a general facility for reasoning about recurrent events, few temporal reasoners have tackled this problem. Describing these will be simplified by first defining a *recurrence* as an event that occurs more than once in a temporal context. A collection of recurrences constitute a *recurrent pattern*.

In one cut at the problem of representing recurrences, all recurrences are enumerated (asserted), as is the interval of the recurrent pattern of which they are members (as is done by Mittal [41]). As a result, the recurrences appear to the temporal reasoner no differently than other, non-recurring events. The interval of the recurrent pattern similarly appears as just another interval with temporal relations arranged so that each recurrence is during that interval. It is hardly feasible, however, to enumerate highly repetitive recurrent patterns with large numbers of recurrences (e.g. the cardiac rhythm).

Kandrashina's [26] "T-Model," partially implemented in the VOSTOK system, follows another approach to the representation of recurrent events. In this approach, several different types of recurrent patterns ("chains" in the T-model representation) are defined, and a set of generic relations between these patterns is specified. As in Figure 2.18(a), the "chains" of the P waves and QRS complexes are synchronized in a particular relationship (ALTERNATION in T-model terminology) with respect to one other during the course of the normal recurring pattern of cardiac electrochemical activity.

Where this last approach performs poorly is in specifying occasional modifications in the temporal relations of the recurrent pattern. Take for example, an

²⁶Notable exceptions include the survival curves for various diseases and compilations such as the Denver Developmental Screening Test [18] which charts the percentile of patients that manifest a particular sign of development at different ages.




Figure 2.18: Patterns of P Waves and QRS Complexes in ALTERNATION.

electronic recording of thousands of cardiac cycles, in which one (Figure 2.18(b)) or more recurrences have prolonged P-R intervals. It is difficult to use a generic recurrence representation, of the T-model type, to represent this information and in particular to make such temporal deductions as:

- The temporal relations between the events on opposite sides of the perturbed recurrence. This requires combining the generic information of the recurrent pattern and that of the local variations of that pattern.
- The total duration of the pattern.

Obviously, a constraint propagation mechanism could make all these inferences, but only if all recurrences were enumerated—hardly an elegant solution and generally not feasible.

What is required is a dichotomous temporal reasoning system. One part would represent, by explicit enumeration, those temporal relations that differ from the recurrent pattern. The other side of the system would represent the generic recurrent patterns. A mapping would have to be maintained, linking the recurrent patterns to the enumerated events. One of the major difficulties in the design of such a system would be to maintain the mapping and spotting inconsistencies between the two halves of the dichotomous representation.

THRIPHT and TUP together implement a poor man's version of the above scheme. TUP maintains the temporal relations of the enumerated portions of the recurrent

patterns. THRIPTH maintains recurrent patterns as abstracted events²⁷ with *ad hoc* specification of the maximum number of references and the bounds on the interval containing the entire recurrent pattern. Clearly, this is not satisfactory, and the implementation of the system, described above, is one of the more pressing items on the agenda for further research in temporal reasoning.

2.9.3 Parameterized Bounds

In a patient history, there rarely is any need, or sufficient medical knowledge, for the use of parameterized bounds. However, when the diagnostic domain is one of the more thoroughly investigated physiologies, such as acid-base homeostasis, there is sufficiently detailed medical knowledge to take advantage of parameterized bounds of the kind illustrated in form 10. As shown, instead of explicit numerical values, the bounds are represented as functions of vectors of quantities.

Form 10

(RREL <point one> <point two> f(Q1) g(Q2))

The implementation of parameterized bounds would require that every time the value of a quantity changed, all RRELS whose bounds depended upon that quantity would have to be updated. If the constraint of these RRELS were to increase, constraint propagation would proceed as previously described. If the change in the quantity caused a *decrease* in constraint in any RRELS, TUP would behave as if it were withdrawing these RRELS.²⁸

The fact that TUP does not represent RRELS as probability distributions creates some difficulty in using parameterized bounds. Take for instance, the use of parameterized bounds to describe such functional dependencies as the expected period of survival of an individual exposed to high levels of radiation, as it varies with the cumulative exposure and the nature of this radiation. Unfortunately, because the temporal estimates that arise from parameterized bounds are usually of a

²⁷Which then become the reference sets of the recurrences.

²⁸That is, recursively withdrawing all those RRELS whose state of constraint depended upon the state of the RRELS whose constraint had been weakened.

probabilistic nature,²⁹ temporal inference errors would be generated, as previously discussed.

²⁹In the case of the radiation exposure, there is a probability distribution of survival for each exposure level.

3. Application of TUP to Medical Diagnosis

THRIPHT is a medical expert system prototype, developed in order to investigate the kinds of temporal reasoning required for the task of medical diagnosis. Early in its design, I made the decision to avoid a relatively simple graft of TUP onto a "vanilla" rule-based system. This would have involved adding TUP predicates to those already available for the construction of rules, and permitting the assertion of TUP temporal relationships in the data base as in example 16.¹ As the limited range of knowledge representation and the limited control structure that such a system provides would have been inadequate to support the scope of the multi-stage diagnostic task outlined in the introductory chapter, I have instead chosen to build a prototype expert system using several of the "second generation" AIM technologies. These include: causal aggregation hierarchies (e.g. ABEL [46] and CADUCEUS [50]), knowledge structures for reasoning about hypothesis evocation distinct from those employed to perform causal reasoning within a hypothesis (e.g. CADUCEUS' constrictor and causal hierarchies), and the concurrent construction of several alternate patient models rather than building and dismissing hypotheses one after another. This system—THRIPHT—is skeletal, incomplete and far more brittle than the systems whose concepts it borrows. It has the merit, however, of bringing together in one system the features that are necessary for the investigation of the temporal issues in medical diagnosis.²

Example 16

```
IF (BEFORE-NOW-P
    ((NAME JAUNDICE)
     (TYPE BEGIN-INTERVAL)))
AND
(BEFORE-NOW-P
```

¹Before the medically knowledgeable reader protests, I emphasize that this example is meant to illustrate the use of TUP predicates and assertions in a rule and not define the necessary criteria for hepatitis.

²Its implementation also had the merit of not exceeding the available (my own) man-power.

```

((NAME IV-DRUG-ABUSE)
 (TYPE BEGIN-INTERVAL)))
AND
(S-BEFORE-P ((NAME IV-DRUG-ABUSE)
             (TYPE BEGIN-INTERVAL))
            ((NAME JAUNDICE)
             (TYPE BEGIN-INTERVAL)))
THEN
(RelationToPresent
 ((NAME HEPATITIS)
 (TYPE BEGIN-INTERVAL))
 (+EPSILON) (+INFINITY))

```

Whereas the previous chapter dealt principally with temporal representation and isolated examples of temporal reasoning in the medical domain, in this chapter the interaction between the temporal reasoning system (TUP) and a medical expert system (THRIPHT) and the role of this interaction in the various phases of diagnosis is emphasized. For this purpose, we return to the scheme (elaborated in chapter 1) for decomposing the task of history-driven diagnosis. I repeat that this is a pragmatic but still rather arbitrary division of labor in the diagnostic process. I do not claim any close analogy to human cognition or that this scheme is any more "correct" than others. It does however conveniently organize the exploration of different types of diagnostic behavior and does follow the general pattern of many "second generation" AIM systems.

3.1 Data Collection

In the course of AIM research, several different paradigms for obtaining patient data have been experimented with. On one extreme, there are inflexible, "dictatorial" systems that specify what information they require, and are unable to perform further diagnostic activity until these requirements are met. On the other extreme, the passive systems accept whatever data is preferred (or is available) and then proceed to make whatever conclusions are possible with this information. In between, the mixed initiative programs accept all information volunteered but may also prompt the user to provide some pertinent details.

In the current implementation of THRIPHT, the data-collection phase is largely of the passive type, whereas during the hypothesis evaluation phase it is of the mixed initiative type. In earlier versions of THRIPHT, the mixed initiative paradigm was also adopted for the data-collection phase so that a few heuristics could be used to ensure the consistency and precision of the data. If externally asserted RREL's were

found to have a CI much higher than average, THRIPHT would phrase a question to attempt to obtain more constrained information. Similarly, THRIPHT would check selected derived RREL's by asking the user if the derived information corresponded to her understanding of the temporal distribution of events.³ THRIPHT would also initiate a "canned" set of questions for significant findings; for instance if melena⁴ was reported, this would cause THRIPHT to query the user about related events such as pallor, loss of weight or anemia. As described earlier, THRIPHT has been designed as if it were to operate in the background of an automated medical record system using the information gathered to check if any important or likely diagnostic hypotheses had been omitted. Therefore, the current version of THRIPHT is purely passive in its data-gathering phase.

All assertions that are made to THRIPHT in this data-gathering phase are asserted as RREL's in the REALITY context—the default when no context is specified. Constraint propagation is performed upon the assertion of each RREL. This permits the frequent temporal inconsistencies in the history to be detected early in the diagnostic process before any significant effort is wasted upon hypotheses based upon erroneous data. As this data-collection occurs prior to the generation of any hypotheses there isn't any applicable domain knowledge to guide the clustering of the temporal data base into reference sets. During this phase, reference sets could only be generated arbitrarily and with a higher risk of missing important inconsistencies. For this reason, all RREL's are asserted within the same reference set to ensure global consistency. As this phase usually involves no more than 20-30 point events,⁵ exhaustive range addition is not too computationally demanding.

Whether implicitly (through the tense of a verb) or explicitly, most temporal assertions obtained in the patient history contain a reference to the temporal position of the specified event(s) with respect to the present. To represent this, the RelationToPresent assertion is used with virtually every datum gathered in this first phase. For instance, the phrase "the onset of fever *began* two days before the rash" not only indicates the relative position of the fever and rash but also the fact that the onset of the fever occurred prior to the present. Example 17 provides the

³For instance if the patient asserted that the fever happened two to three days before the rash and the rash four days before the jaundice, the patient could be asked if she believed that the fever preceded the jaundice by six to seven days.

⁴Stool stained by blood pigments.

⁵Experienced physicians will generate a small group of working hypotheses [27] early in the taking of a patient history.

equivalent assertions in TUPese.

Example 17

```
(RREL ((NAME FEVER) (TYPE BEGIN-INTERVAL))
      ((NAME RASH) (TYPE BEGIN-INTERVAL))
      (2 DAYS) (2 DAYS))
```

```
(RELATIONTOPRESENT ((NAME FEVER) (TYPE BEGIN-INTERVAL))
                   (+EPSILON) (+INFINITY))
```

The consequence of the ubiquitous reference to the present is that all events in the history become solidly anchored with respect to the present and indirectly (by means of the expansion of the RelationToPresent assertion—see example 6 on page 24) to the calendar reference system. As the calendar reference system has a very low overall constraint index, these frequent references to the present are helpful in constraining the temporal relations of the patient history.

3.2 Hypothesis Evocation

There usually are many hypotheses that include findings that correspond to those in the patient history and yet which a human expert would only briefly consider, if at all. One reason for this is that for a large class of hypotheses, their temporal patterns are wildly at variance with the chronology of the events in the patient history. Therefore, to correctly trigger hypotheses for further consideration, the expert system must be sensitive to the temporal configuration of events in the patient history. Unsatisfactory results will ensue if just the presence or absence of events is used to evoke a hypothesis. This is clearly illustrated by the three examples of mistaken diagnosis of transfusion-borne acute hepatitis B in the introductory chapter.

Temporal Errors In Human Expertise

Human beings who ignore temporal course are subject to the same errors and inefficiencies to which temporally unsophisticated expert systems are prone. One such error will have occurred to anyone who has taken several patient histories. A patient will mention several symptoms, say a prior episode of jaundice and malaise,

and is then asked routine questions to refine the diagnosis—in this case whether the patient has received blood transfusions. An affirmative answer will occasionally incorrectly lead the interviewer to believe that she is on the right track only to discover that the event happened at a time other than one relevant to the hypothesis (e.g. the transfusion occurred *after* the jaundice). That is, the hypotheses the interviewer considers may lead her to expectations that, without sufficient temporal cues, will be incorrectly matched to findings reported by the patient. The converse phenomenon will also occur: the patient omits a symptom because it is outside the temporal window the patient considers to be relevant. A classical case of this is the patient who has been diagnosed as having had a myocardial infarction or end-stage emphysema and is asked if he smokes. Often the patient will respond in the negative, and it only becomes apparent, later in the interview, that the patient started to abstain a month previously, after a fifty-year, two-pack-a-day history of smoking.

The Usual Approach

When a medical expert system is developed, the system's inability to recognize a wide variety of temporal patterns of disease leads the knowledge engineer to assume, often implicitly, one or both of the following two "solutions".

The first solution is to assume that the user will go some distance in arriving at a diagnosis and therefore only "feed" the expert system those events that are pertinent to the present illness. In a complete patient history that includes multiple visits⁶ such assumptions can lead to erroneous conclusions. For example, MYCIN tests for a bacteroides infection with the following rule [12]:

If (1) the infection is primary-bacteremia, and
 (2) the site of culture is one of the sterile sites, and
 (3) the suspected portal of entry of the organism is the
 gastrointestinal tract,
 then there is suggestive evidence (.7) that the identity of the
 organism is bacteroides.

Regardless of when the primary bacteremia occurs it can bind to the corresponding slot in the rule. Therefore, even if the primary-bacteremia happened months before the current culture was obtained, MYCIN would come to the same conclusion

⁶ After all, one of the goals of AIM is to produce programs that operate over the whole lifetime of a patient and that take into account events of clinical importance throughout.

as if the bacteremia had happened just a few hours before the culture. Also, if the patient had multiple episodes of primary bacteremia, each of the bacteremia events could be matched to the antecedent part of the rule. The user is therefore expected to assert bacteremia only if the episode was recent enough to be relevant to the present illness. In doing so, the user has to engage in diagnostic activity because temporal proximity does not necessarily imply relevance to the present illness.⁷ For instance, an episode of acute rheumatic fever that causes valvular disease 30 years later should be reported to the expert system even if it is not temporally proximate to the present illness. If we intend to develop AIM systems that can work with naive users we cannot shift it any of the responsibility for recognizing relevant findings from the expert system.

An alternate solution, and probably the one most commonly employed, is to permit the program to trigger freely upon the whole patient history and let the expert system discover⁸ at a later phase that the facts do not quite fit. For instance, if the post-transfusion hepatitis B hypothesis was triggered by the jaundice that had occurred in infancy and a recent transfusion during caesarian section, the fact that the serologies (for the various hepatitis B antigens and antibodies) were all negative would most likely cause the hepatitis hypothesis to be ranked low, if at all, in the differential diagnosis. This method, however, involves the exploration of many solutions (and requests for information) that are irrelevant at a glance to clinicians familiar with the disease chronology. Often, even if all the atemporal information is available and is used, the expert system will be unable to make distinctions for which temporal representation is necessary. It is apparent, then, that providing expert systems with the capability to represent and reason about the temporal course of disease permits a significant pruning of the diagnostic search space—even at the hypothesis triggering phase.

Consequently, if a large-scale diagnostic program is endowed with the ability to

⁷This criticism of MYCIN may appear unfair since the system "backward-chains" from the hypotheses to the data and therefore the timing of the events can be explicitly checked by querying the user. There are several flaws with this approach: (1) the user or data base must provide the specific temporal relationship asked for; MYCIN cannot infer the temporal relationship from prior assertions. (2) in this form of backward-chaining, MYCIN would exhaustively ask questions of the timing of events in each hypothesis considered. This is unacceptable if a patient history is to be taken from a patient whose underlying problem is unknown—backward-chaining would then have to be initiated from too many hypotheses. (3) if the rule is used in a forward-chaining mode, my original claim, that the user has to take an active role in generating the differential diagnosis, holds.

⁸On purely atemporal grounds.

support sophisticated temporal reasoning, its diagnostic style shows a greater degree of focus; that is, a more *relevant* set and *smaller* number of diagnostic hypotheses are considered.

How Much is Gained?

Let us put aside, for the moment, the reduction in the number of questions that a more focused diagnostic style would bring to ask the following question: To what extent does the use of temporally sophisticated hypothesis evocation lead to improved expert system performance?

One way I could show that the computational savings are substantial would be to pit THRIPT against an expert system without the capability for temporal reasoning. Both systems would be presented with a set of run-of-the-mill cases that a physician might see in the course of a week's practice. The time that each expert system would require to reach a final diagnosis on this typical case-load would provide statistical evidence of the relative performance of the two systems. This method has not been tried because THRIPT lacks a broad knowledge base and, as mentioned before, the ranking (fifth) phase of diagnosis is only partially implemented. Nonetheless, the scope of the reduction in computation can be conveyed by considering the three following sets of high-prevalence findings.

1. Take the findings of a non-specific rash and bee-sting. Very many patient histories will contain both of these events. There are many causes of a rash but only a small percentage of these will be caused by a bee-sting. However, if a medical expert system is unable to check whether the rash followed the bee-sting by a few minutes, it will have to give further consideration to the hypothesis of an allergic reaction secondary to a bee-sting in most, if not all, histories which contain both findings.
2. Many women will lactate during the course of their lives. The vast majority will do so towards the end of their pregnancy and during nursing. A physician will not consider pathological causes of lactation, such as thyroid disease or a pituitary adenoma, if the lactation occurs during these periods when it is expected.⁹ However, a temporally incompetent expert system will have to

⁹Unless the diseases have progressed to the point where they are clinically manifest in ways other than lactation.

further investigate all these cases thereby using additional computational and clinical resources.

3. Many patients who are administered antihypertensive drugs will have rebound rises in their blood pressure if they are weaned too quickly from their medication. This will occur when the therapy is either terminated or a different drug is selected. The new medication might be given in insufficient dosage or lacks efficacy for that patient. As there are many other causes for increased blood pressure, unless the change in medication precedes the rise in blood pressure by a specified period, it is a waste of computational and clinical resources to devote more effort to a hypothesis based upon changes in antihypertensive medication.

In general, many findings which are common in the general population are also drug side-effects. Each drug has a characteristic delay between its administration and the effects/side-effects. If an expert system is able to represent this delay and then match it against the timing in the patient history¹⁰ then the hypotheses of drug-induced effects need only be considered in a small fraction of the prevalent cases of the findings.

These examples illustrate how historical items lose their specificity if they are stripped of the description of their place in a characteristic chronology. At least in medicine, precedence information alone is insufficient to adequately characterize such chronologies and the quantitative timing information imparted by RREs is required.

Early Solutions

Developers of early AIM programs were cognizant of the need for some sort of temporal representation for triggering hypotheses. The *Present Illness Program* (PIP) [48] for instance was modified [55] so that features of hypothesis frames such as causal and associational links could have a temporal qualifier. To this end, the time-line was divided into five periods: PAST, RECENT-PAST, NOW, NEAR-FUTURE and FUTURE. This enabled PIP to represent disease chronologies such as the development of chronic glomerulonephritis (CGN) in the FUTURE from acute glomerulonephritis (AGN) NOW. PIP could then automatically infer when features characteristic of

¹⁰ And, if necessary, derive the necessary timing information from prior temporal assertions in the patient history.

a hypothesis could be expected with respect to the present. Consequently, PIP would not ask the user questions about features occurring during CGN if AGN was hypothesized to occur NOW.

This *ad hoc* solution has limitations: a large number of disease states have to be created,¹¹ calculations of relative temporal position lead to loss of temporal precision,¹² and quantitative temporal information, as opposed to simple ordering information, cannot be represented or reasoned with in this manner. Other approaches have employed the precedence relations between causal antecedents and consequents but, with branches in the causal chains, many inter-event temporal relations are indeterminate.

THRIPHT's approach

THRIPHT triggers take full advantage of TUP's richness of temporal expression. A THRIPHT trigger can consist of any boolean combination of predicates, temporal or atemporal. Hypotheses are triggered only if characteristic collections of signs and symptoms are present and only if they are in a temporal configuration compatible with the hypothesized disease. Example 18 illustrates part of the trigger for the Hepatitis B hypothesis. This example roughly corresponds to "consider acute hepatitis B secondary to intravenous drug abuse if (in addition to the other criteria not shown) jaundice has been manifested, intravenous drug abuse did not end before the previous seven months, and the onset of drug abuse preceded the onset of jaundice by at least two months."¹³ Observe that all predicates (by default) test the REALITY context. Also, the trigger in example 18 has been written with the relaxed version of the TUP predicates so as to make the trigger more sensitive (and less specific).

Those TUP queries that employ the RelationToPresent retrieval form (see section A.1.4) generate, as side-effects, additional assertions. So these will not clutter up the patient history, a query context is created¹⁴ for these side-effect assertions.

¹¹Especially if temporal distinctions, finer than the five periods provided, are to be made.

¹²e.g. If a clinical feature *X* is asserted to occur in the NEAR-FUTURE with respect to the AGN hypothesis, and if AGN is asserted to have happened in the RECENT-PAST, PIP could not determine if *X* would occur NOW or in the NEAR-FUTURE.

¹³The two and seven months refer to the minimum and maximum incubation times, respectively, of the virus.

¹⁴THRIPHT simply adds the context specification to the new points.

The query context shares with the REALITY context the assertions of the patient history. After the queries are completed the side-effects can be eliminated by garbage-collecting the contents of the query context.

Example 18

```
(AND
  (BEFORE-NOW-BY-P
    ((NAME JAUNDICE) (TYPE BEGIN-INTERVAL))
    (+EPSILON) (+INFINITY))
  (BEFORE-NOW-BY-P
    ((NAME IV-DRUG-ABUSE) (TYPE END-INTERVAL))
    (-INFINITY) (+7 MONTHS))
  (BEFORE-BY-P
    ((NAME IV-DRUG-ABUSE) (TYPE BEGIN-INTERVAL))
    ((NAME JAUNDICE) (TYPE BEGIN-INTERVAL))
    (+2 MONTHS) (+INFINITY)))
```

Unstructured collections of triggers will tend to cause far too many hypothesis triggers to be tested and then too many of these to be successfully activated. Several schemes have been elaborated, the most successful of which has been to create disease classification hierarchies so that only the relevant disease categories are even considered for triggering. The scheme adopted for THRIPT is a variant of the one used in MDX [10] in which there is a community of *specialist* programs organized in hierarchical fashion. In the MDX system, each specialist attempts to establish or refine a diagnosis. THRIPT instead uses the hierarchy to shrink the number of hypothesis triggered and then considered for further diagnostic activity.

In general, each node in the hierarchy represents a category of pathophysiological pathways (diseases) which should be considered if certain conditions (the trigger) are satisfied. Each node really has three associated triggers: RULE-OUT, NECESSARY and SUFFICIENT. If the conditions of the RULE-OUT trigger are satisfied, the hypothesis is excluded from further consideration. In the current implementation, triggers are categorical links between findings and hypotheses so that satisfaction of either of the NECESSARY and SUFFICIENT triggers causes the active consideration of the hypotheses with which they are associated.

Evaluation of triggers proceeds from the broadest categories to the more specific. Every time the trigger conditions of a node (pathophysiological hypothesis) in the

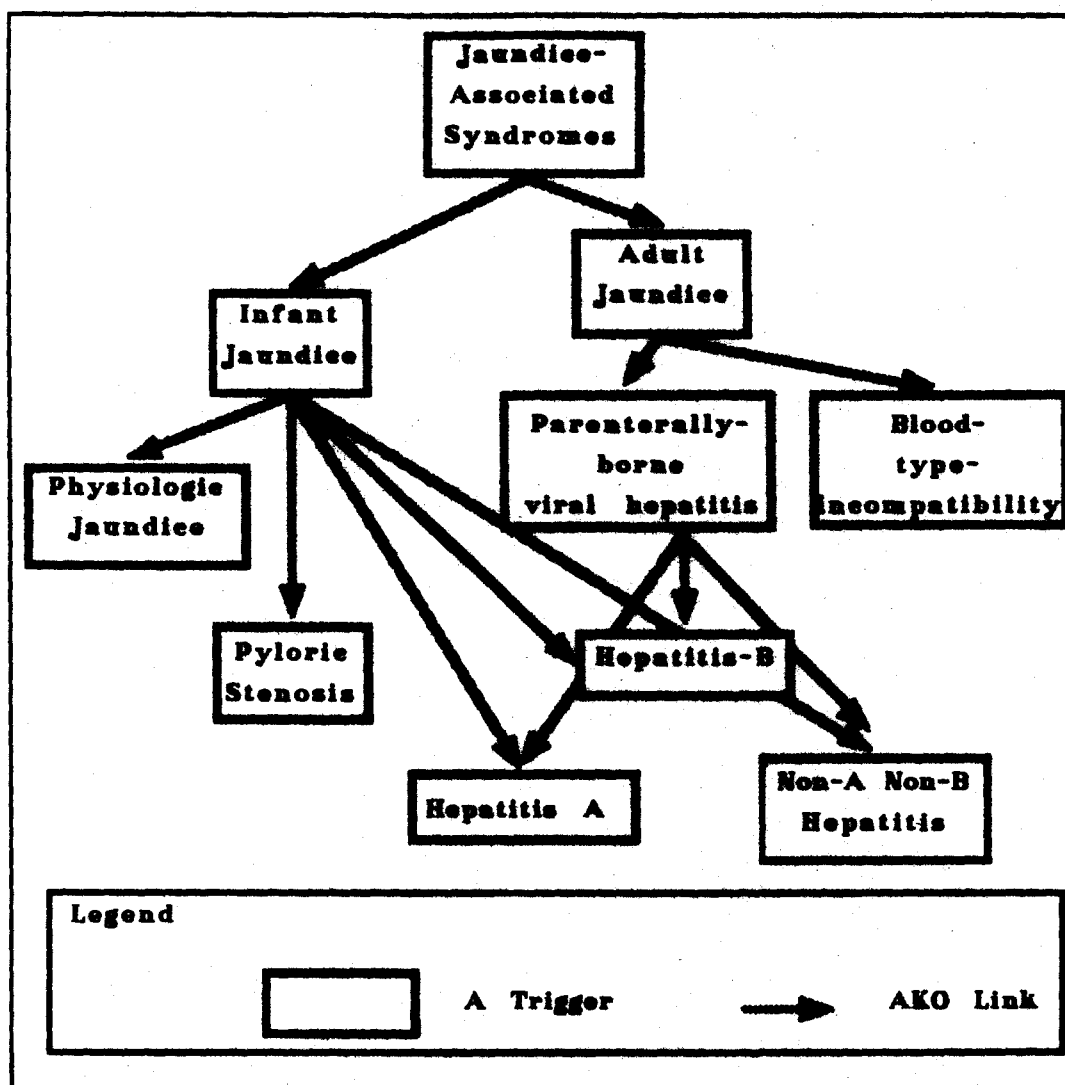


Figure 3.1: Part of the Tangled Hierarchy of Triggers.

tangled hierarchy are satisfied, all the triggers of the specialized instances¹⁵ of that node are evaluated to ascertain whether one or more of these narrower hypotheses should be considered. This is repeated recursively until the most specialized nodes have been triggered. Once this is achieved, only the pathophysiological hypotheses at the most specialized triggered nodes are processed in the next—hypothesis elaboration—phase.

Regarding the sample tangled hierarchy in figure 3.1, note that THRIPHT does not compel the user to enter the jaundice as ADULT-JAUNDICE or INFANT-JAUNDICE as many systems do.¹⁶ Rather than forcing the generation of many compound terms, THRIPHT enables the knowledge engineer to employ TUP predicates to test for the temporal relation of findings (e.g. jaundice to age). These temporal relations may have been directly asserted in the medical record or indirectly obtained from other temporal relations by means of TUP's inference mechanism.

3.3 Elaboration

Hypotheses that are triggered in the evocation phase are in fact collections of pathophysiological pathways (see figure 2.15, page 47). What brings together these pathophysiological pathways depends upon the classification principle used in the trigger hierarchy. It may be that these pathways share clinical manifestations, share etiology and/or share pathophysiological mechanisms. The purpose of this, the elaboration phase, and the phases that follow, is to manipulate these pathophysiological subhypotheses so that they may be distinguished from one another. The particular role of the elaboration phase is to separate each of these subhypotheses, as diagrammed in figure 3.3.

Some of these subhypotheses are compatible with one another. In the hepatitis B hypothesis there are subhypotheses with jaundice and with an immune response, each of which can occur by themselves or together. Other courses, such as immediate convalescence after the acute phase of hepatitis, and chronic active hepatitis are by definition mutually exclusive. During the elaboration phase, the compatible pathways (subhypotheses) are grouped together in classes of compatible hypotheses whereas mutually exclusive subhypotheses are segregated into competing classes of

¹⁵n.b. Daughter nodes trigger subsets of the diseases their parent(s) triggers. The daughter nodes are therefore linked to their parent(s) by AKO links.

¹⁶In those cases when at least some *ad hoc* attempt is made to include temporal information.

hypotheses.

3.3.1 Defining Mutually Exclusive Courses

The exercise of representing the course of disease as an entity that progressively changes over time makes it very clear that few diseases are truly mutually exclusive. A program that does not represent disease chronology might represent metabolic acidosis and metabolic alkalosis as mutually exclusive pathophysiological disease states. However, a patient can have both, as in the case of the patient suffering from aspirin intoxication, because the mutually exclusive events are not cotemporal. One of the criteria for the mutual exclusivity of hypotheses is therefore the cotemporality of mutually exclusive events. Dean's [14] temporal reasoner is designed towards detecting just such cotemporality. In practice however, determining cotemporality of events in independent hypotheses is difficult. One problem is that each hypothesis has to be anchored with respect to a shared point in time, such as the present, which the hypotheses are not until instantiated as patient models. Therefore it is *a priori* almost¹⁷ impossible to rule-out the co-existence of two diseases, in one patient, just because the two diseases contain mutually exclusive events. The only information that we can obtain *a priori* is that some diseases might only very rarely occur in the same patient. Obviously this does not have the same strength as asserting logical contradiction.

Even when two hypotheses are anchored with respect to the present, it is not easy to determine if logically contradictory events are cotemporal unless the events are tightly constrained. Otherwise, especially if the contradictory events are of brief duration, the temporal information may be insufficiently constrained to determine whether the two events are strictly cotemporal.

Even if it can be determined that contradictory events are cotemporal, there remains the problem of keeping the temporal reasoner free of domain-dependent reasoning. In the general case, it will not work to simply permit the most recent assertion to "clip" the duration of the earlier, opposing assertion. What is needed is a domain-dependent means of evaluating the relative belief in, or likelihood of, the conflicting events.

The heuristic solution implemented in THIRPHT is one that is not satisfactory in the general case. That is, those event pairs that the knowledge engineer knows

¹⁷Almost, because some diseases by definition cannot occur in the same lifetime as another disease, as in the case of systemic lupus erythematosus and rheumatoid arthritis.

to be always contradictory, or for which the probability of occurring in the same patient is known to be very small, are labeled as mutually exclusive in the THIRPHT knowledge base. THIRPHT traverses the causal aggregation hierarchy to ensure that all events in this hierarchy below the two explicitly contradictory events are also marked as contradictory. It also labels all the causal consequents of each event to be mutually exclusive. Even though this is an *ad hoc* solution, it works in most cases. Take, for example, two of the possible sequelae to Hepatitis B: the chronic active state and the chronic carrier state. Although it is conceivable that in a single lifetime, a patient might have both (but at different times), it might be a waste of computational resources to consider this possibility.

A more principled solution would be to prune those combinations of hypotheses calculated to be of low probability. However, as THIRPHT is not equipped with the capability to represent and reason probabilistically about events, this solution is not available. Even if THIRPHT did have such a capability, it is doubtful whether the probabilities for all the disease combinations would be available (see Szolovits and Pauker [56]).

3.3.2 Generating Contexts

For each class of compatible hypotheses, there is an associated temporal context. This context contains the default temporal constraints upon the component events of these hypotheses. These constraints are default ones in the sense that they represent the bounds on the temporal relations for the entire subpopulation of patients described by a particular hypothesis. In the case of any individual patient, these bounds are usually further constrained.

THIRPHT's Default Temporal Assertions

Other than the temporal relationships entered by the knowledge engineer, THIRPHT automatically asserts those temporal relationships that are obvious. For instance, for each event in the knowledge base, a corresponding interval is asserted¹⁸ Also, for all causally connected events, the onset of the cause is asserted to precede the onset of the effect. Similarly, the onset of each manifestation is asserted to be no earlier than the onset of the pathophysiological event that it manifests. The clinical

¹⁸Using the ASSERT-INTERVAL form and therefore with default bounds of 0, +∞ between onset and end of the interval.

manifestation of a causal antecedent many occasionally follow the manifestation of a causal consequent and so prior assertions are not automatically made about temporal relationships between manifestations.

Constraints Set by the Knowledge Engineer

THRIPHT permits the knowledge engineer to bind any number of temporal constraints to each event and between any two events in a hypothesis. These constraints can be any assertional form known to TUP. Upon elaboration of a hypothesis, all these temporal constraints are asserted so long as they constrain events within the same compatible hypothesis class. THRIPHT appends the hypothesis class object to each temporal assertion made so that the hypothesis class becomes the temporal context of these assertions. As a result, each hypothesis class becomes paired with a temporal context. This pairing demonstrates the separation between the temporal and atemporal parts of each pathophysiological sub-hypothesis.

3.3.3 Generating Reference Sets

During the generation of temporal contexts, THRIPHT provides the information that TUP uses to organize the events into reference sets. In the implementation, the REFSET specification is appended to each point specified in each RREL prior to assertion in a hypothesis context. The constraint propagation that follows is then restricted to these reference sets.

As discussed earlier, the causal aggregation hierarchy is used to determine reference set membership. First, the value of the NAME descriptor in a temporal assertion is used to obtain the corresponding (event) node in the THRIPHT causal network. All that is then required to obtain the reference set identification is to find the node's immediate causal aggregation. If the node belongs to more than one causal aggregation, the corresponding time points will have multiple reference set memberships.

3.4 Instantiation

The hypotheses elaborated from the composite hypotheses (that were triggered by the patient findings) are not patient specific. Rather they represent shared characteristics of patient populations with the hypothesized disease courses. It is only when these hypotheses are modified with the information that was obtained from

the patient (history) that the hypotheses become patient-specific. A hypothesis modified with such patient information becomes a patient (specific) model.

3.4.1 Consequences of Constraining a Hypothesis Context

In the creation of a patient model, TERIPHT adds the RREL's gathered in the patient history to those of the temporal contexts paired to each hypothesis (hypothesis context). TUP responds as usual to these additions (assertions) by propagating the constraints through each of these contexts. The global constraint index of a hypothesis context will be significantly decreased since the temporal relationships between events in the patient history will be generally far more constrained than the corresponding relationships in the hypothesis contexts.

Another consequence of constraint propagation during instantiation is that the hypothesis contexts become anchored with respect to the present. All that is required is that there be at least one RREL in the patient history that relates an event to the present or to the calendar.¹⁰

To illustrate, take but one RREL from the Hepatitis B chronic-active hypothesis as in example 19. If the patient history includes the information of example 20, upon instantiation the patient model will include the information that the inoculation occurred four weeks ago and the onset of symptoms 50 to 180 days later (that TUP will deduce to be approximately three to twenty-two weeks in the future).

Example 19

```
(RREL ((NAME INOCULATION) (TYPE BEGIN-INTERVAL))
      ((NAME SYMPTOMATIC) (TYPE BEGIN-INTERVAL))
      (50 DAYS) (180 DAYS)
      CHRONIC-ACTIVE)
```

Example 20

```
(RelationToPresent ((NAME INOCULATION) (TYPE BEGIN-INTERVAL))
                   (4 WEEKS) (4 WEEKS)
                   REALITY)
```

¹⁰TUP uses the real-time clock of the host system, by employing the CALENDAR reference system (see section 2.1.2) to infer the position of points on the calendar with respect to the present.

One important consequence of the constraint propagation during instantiation is that the constraints of the patient history may be found by TUP to be inconsistent with those of the hypothesis context. Thus, even before the hypothesis evaluation phase, some patient models may be dismissed because of temporal inconsistency. In the last example, if the patient history were also to contain the information of example 21, this would place the onset of symptoms within four weeks of inoculation which is inconsistent with the minimum time specified in the hypothesis context (example 19). TUP's constraint propagation mechanism would discover this inconsistency upon assertion of the inconsistent RREL. This illustrates yet again that providing medical expert systems with a sophisticated temporal utility can greatly reduce the problem space that has to be searched—in this case by immediately dismissing temporally inconsistent patient models. An AIM system not so equipped would fail to recognise that the patient data of this last example was inconsistent with the Hepatitis B hypothesis and would therefore go on to expend needless computational effort in attempting to confirm or rule it out.

Example 21

```
(RelationToPresent
  ((NAME SYMPTOMATIC)
   (TYPE BEGIN-INTERVAL))
  (2 DAYS) (3 DAYS)
  REALITY)
```

Multiple Findings, Multiple Event Instances

The process of discarding temporally inconsistent patient models goes a long way to solving the problem of binding multiple instances of findings in the patient history to multiple occurrences of the same finding in a hypothesis context. To wit: in chronic-active hepatitis there may be a febrile episode during the acute period after inoculation and several episodes later during the chronic active phase. The patient data may also contain reports of one or more episodes of fever. A medical expert system not equipped with a temporal reasoner has to create a patient model for each of the possible combinations of bindings of findings from the patient history to events in the disease hypotheses. By employing a temporal reasoner such as

TUP, discrepancies between the patient history and the hypothesis context²⁰ are all detected by the consistency check in the constraint propagation process. THIRPHT is therefore able to discard all those binding configurations that are *a priori* temporally untenable.

3.5 Ranking Patient Models

Ranking hypotheses in THIRPHT is roughly patterned after the examples of ABEL and CADUCEUS, in that the process can be viewed as a repeating loop which has three components: patient model *evaluation*, finding features of the patient models that makes them distinct from others—*differentiation*—and obtaining patient *information* to exploit the distinguishing features. THIRPHT's use of temporal information in this phase, because it accentuates the differences between patient models, causes a more rapid convergence to a "correct" differential diagnosis. It also qualitatively changes the diagnostic style of the expert system, as described in the following sections.

Unlike the four other phases described in this chapter, this phase has not been fully implemented—the main issues involved in such an effort lie outside the scope of the present research effort.

3.5.1 Evaluating Hypotheses

To evaluate a set of hypotheses is to assign a partial order to the members of this set according to the degree of belief in the extent to which each hypothesis accurately reflects reality—in this case the patient's state of health. In the absence of temporal reasoning, medical expert systems focus on scoring hypotheses based on the findings that match events in the hypotheses and those unaccounted for in the hypotheses. TUP, in addition, enables an evaluation of temporal consistency of each hypothesis with the patient data every time the evaluation-differentiation-information loop is performed.

Because the range relation is devoid of probabilistic information, the global constraint index of each hypothesis context cannot be also used by a rank scoring function. That one patient model should have a more constrained hypothesis context than another does not necessarily signify a better fit of patient data to hy-

²⁰In event ordering, duration or position with respect to the present.

pothesized pathophysiology. Therefore, other than the consistency check, temporal information, as it is represented in TUP provides little help in calculating a rank of a hypothesis. Note this does not prevent ranking the hypothesis on atemporal features, as other AIM programs have done.

Information Gathering

Temporal "common sense" in information gathering is important. A medical expert system does not inspire confidence when it requests information obtainable only in the distant future. Nor does it make sense for such a system to repeatedly check for the presence of an event that (to a human) is definitely in the past. THRIPTH has the potential to avoid this sort of blunder because the hypothesis contexts of the patient models are almost always anchored to the present. Consequently, when some datum is to be obtained to distinguish one patient model from another, THRIPTH can determine whether the information about the corresponding event is in the past, straddles the present or is in the future. By judicious use of a few heuristics, this knowledge can cause THRIPTH to exhibit temporal "common sense" in its information seeking behavior. For example:

- Do not consider (not just defer) seeking information about events more distant than m seconds in the future.
- Defer seeking information about events more distant than n but closer than m seconds in the future.
- If an event has a temporal relation to the present that has a lower bound that places the event in the past and an upper bound that places it in the future, be persistent about seeking information about this event.
- Do not aggressively and repetitively seek information about past events that are reported to be unknown.

In diagnosing a jaundiced neonate, THRIPTH could then ask if the serum bilirubin had fallen but would not ask if the patient's liver had yet attained adult function.

3.5.2 Differentiating Between Patient Models

The differences found between patient models drives the diagnostic process in this phase. These differences determine which patient data would be most effective in further separating the scores of the patient models in the differential diagnosis.

TUP's temporal representation enables THRIPT to distinguish between hypotheses that differ in the order of events. For instance, the (traumatic) limb amputation that leads to infection and the infection that leads to (therapeutic) limb amputation. Distinctions made between hypotheses that show the same ordering of events but with quantitatively different temporal distribution are also frequently useful. In neonatal jaundice, if the neonate has *breast milk jaundice*, the serum bilirubin (pigment responsible for jaundice) concentration rises progressively from the fourth day of life and reaches a maximum by 15 days of life. Jaundice due to maternal-infant Rh incompatibility is clinically manifest in the first day of life. *Physiologic jaundice* occurs after the first day of life but ends within a week of birth (in the full term infant). Temporal representations such as TUP's can capture these differences in clinical course and expert systems such as THRIPT can locate them, thereby improving diagnostic acuity.

3.6 Annotated Example

In the preceding discussion, I have described the mechanisms of operation of the individual pieces of the diagnostic engine. In contrast, the example below is meant to provide a feel for how these various pieces fit together in the process of diagnosing a patient.

This annotated transcript describes user entries and THRIPT's response. Since, during the course of my research, I purposely did not touch upon the issues of temporal expression in natural language, THRIPT lacks a "user-friendly" interface. Therefore, the English rendering I have given of the dialog with THRIPT, corresponds to Lisp-forms intelligible to TUP and THRIPT. Where appropriate, I show the forms used and diagram some of the structures generated by THRIPT.

The example is a fictitious episode in which a patient's history is taken at 8 a.m., July 5th 1986. Items from the patient history are rendered in a sans serif type style, THRIPT's responses in bold type style and my comments in *italics*.

Beginning of the data-gathering phase

0. The patient was 45 years old on July 1st, 1986.

Equivalently:

```
(RREL ((REFSYSFORM '45 YEARS OLD')
      (REFSYS AGE)
      (TYPE POINT))
```

((REFSYSFORM "00:01 a.m., July 1st, 1986")
 (REFSYS CALENDAR)
 (TYPE POINT)
 (0 SECONDS) (0 SECONDS))

This assertion permits TUP to respond to queries about the patient's age now, or at any other time, and with respect to any other event.

1. The patient received a blood transfusion 9 to 12 weeks ago

Equivalently:

(RelationToPresent ((NAME BLOOD-TRANSFUSION)
 (TYPE BEGIN-INTERVAL))
 (9 WEEKS) (12 WEEKS))

The expansion of RelationToPresent is given on page 24. The (DATE) form in the expansion would evaluate to the current date and time: "8 a.m., July 5th 1986"

2. Onset of jaundice occurred 8 to 9 weeks after the transfusion.

Equivalently:

(RREL ((NAME BLOOD-TRANSFUSION)
 (TYPE BEGIN-INTERVAL))
 ((NAME JAUNDICE)
 (TYPE BEGIN-INTERVAL))
 (8 WEEKS) (9 WEEKS))

3. The onset of jaundice was first manifest in the morning, between 8 and 10 o'clock on June 5th, 1986

Equivalently:

(RREL ((REFSYSFORM "8 a.m., June 5th, 1986")
 (REFSYS CALENDAR))
 (TYPE POINT)
 ((NAME JAUNDICE)
 (TYPE BEGIN-INTERVAL))
 (0 SECONDS) (2 HOURS))

Redundant data are proferred.

4. Inconsistent temporal information: statements 1., 2., and 3. are inconsistent. Please withdraw one or more of the statements before proceeding.

TUP detects that the bounds are inconsistent. Statements 1. and 2. would have the jaundices occurring as early as June 7th, and as recently as the present (July 5th, 1986). This is inconsistent with statement 3.

5. Statement 3. is incorrect, the onset of jaundice was first manifest at 8 o'clock, in the morning, on June 28th, 1986

The RREL of statement 3. is withdrawn and the new RREL is asserted with the result shown in Figure 3.2.²¹

With this assertion the data-gathering phase ends, and the hypothesis evocation phase begins. A paraphrase of THIRPHT's reasoning is only given when a condition of a trigger is satisfied.

6. As the onset of jaundice has occurred, THIRPHT is considering jaundice-associated syndromes.

The topmost node in the THIRPHT's trigger hierarchy, shown on page 78, has a SUFFICIENT condition that is satisfied by the patient data:

```
(BEFORE-NOW-P
  ((NAME JAUNDICE)
   (TYPE BEGIN-INTERVAL)))
```

This condition brings up for consideration the next most general nodes whose triggers assume that the patient history includes the onset of an episode of jaundice. This represents the effect of my decision that the isolated assertion of jaundice in the future (i.e. speculation) will not cause active consideration of jaundice-associated syndromes.

7. As the age of the patient is greater than one year THIRPHT is considering adult jaundice and has ruled out infant jaundice.

The adult jaundice trigger has a SUFFICIENT condition that is satisfied:

```
(AND
  (RREL ((REFSYSFORM "1 year old")
        (REFSYS AGE)
        (TYPE POINT))
        ((NAME DUMMY-POINT))
```

²¹In the disease hypothesis for hepatitis B, the parenteral introduction of the infectious agent is represented by INOCULATION rather than a BLOOD-TRANSFUSION. The transfusion events are consequently substituted with inoculation events. Although this synonym matching could be done automatically, in the current implementation, I have to provide the substitution.

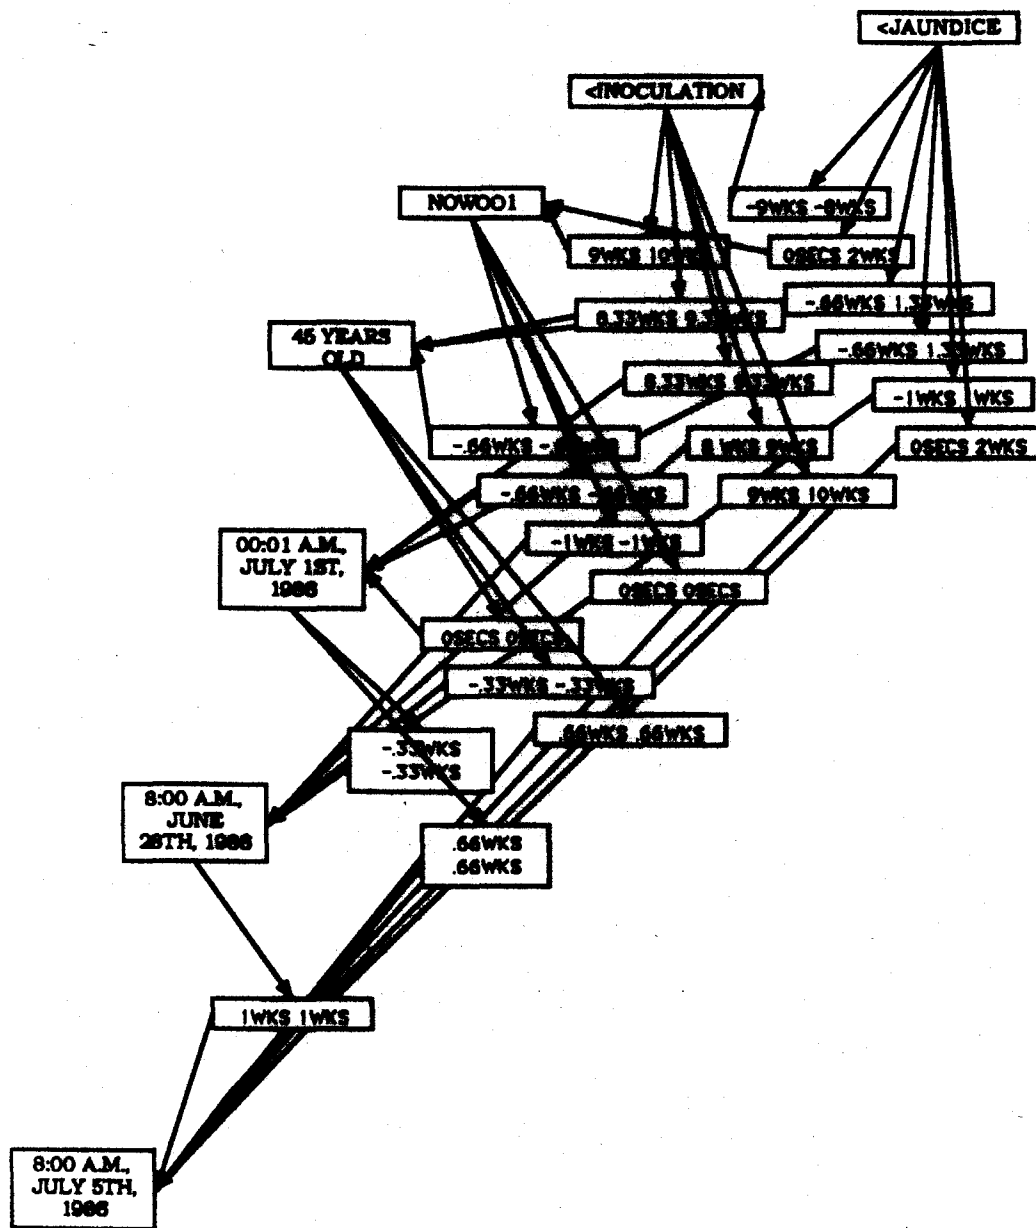


Figure 3.2: The Patient History

```

      (TYPE POINT))
    (0 SECONDS) (0 SECONDS))
  (S-BEFORE-NOW-P
   ((REFSYSFORM "1 year old")
    (REFSYS AGE)
    (TYPE POINT))))

```

The first assertion is made to intern the 1 year old point so that TUP may then perform the constraint propagation that will link that age to the calendar. Recall that the assertions generated during THIRPHT queries are in the query context and can easily be flushed.

The infant jaundice trigger has a RULE-OUT condition that is identical to that of the SUFFICIENT condition of adult jaundices save for the substitution of S-AFTER-NOW for S-BEFORE-NOW.

8. As the transfusion preceded the jaundice, THIRPHT is considering PARENTERALLY-BORNE VIRAL HEPATITIS. As time between the onset of jaundice and transfusion is greater than two weeks, THIRPHT is ruling out BLOOD-TYPE-INCOMPATIBILITY.

The trigger for viral hepatitis caused by parenteral²³ inoculation with the virus has, as a sufficient condition, a check that the transfusion precedes jaundice. Although instances have been reported, hepatitis A infections infrequently occur as a result of parenteral inoculation; the fecal-oral route is more common. Nonetheless, just as I have omitted many hypotheses that should be considered, I have included hepatitis A for demonstration.

The blood-type incompatibility trigger has the following RULE-OUT condition:

```

(BEFORE-BY-P
 ((NAME TRANSFUSION) (TYPE BEGIN-INTERVAL))
 ((NAME JAUNDICE) (TYPE BEGIN-INTERVAL))
 (2 WEEKS) (+INFINITY))

```

Note that the above condition does not test the relation to the present since this has been done by the JAUNDICE-ASSOCIATED SYNDROMES trigger.

9. As the transfusion preceded the jaundice by more than two weeks and less than 120 days, THIRPHT is considering NON-A, NON-B HEPATITIS. As the transfusion preceded the jaundice by more than 50 days and less than 180 days, THIRPHT is considering HEPATITIS TYPE B. As the

²³By means other than ingestion, usually intramuscular or intravenous injection.

transfusion preceded the jaundice by more than 45 days, THRIPT has ruled-out HEPATITIS A.

The three triggers have sufficient and rule-out conditions that test for the respective delays between the onset of jaundice and the transfusion. Observe that it is only the hypothesis of parenterally inoculated hepatitis A that has been ruled out, not all etiologies of hepatitis A. If this example were to be complete, I would show how another hepatitis A hypothesis (with a different inoculation mechanism) was triggered through another path down the trigger hierarchy.

As the hepatitis B and non-A, non-B hepatitis triggers are terminal nodes in the trigger hierarchy, the hypotheses associated with these triggers are evoked and the elaboration phase begins.²³

10. HEPATITIS TYPE B has been elaborated into three mutually exclusive subhypotheses. NON-A, NON-B HEPATITIS has been elaborated into three mutually exclusive subhypotheses.

As illustrated in figure 2.15 (page 47), the hepatitis B hypothesis has three mutually contradictory courses. The first course H_1 goes from acute disease to resolution. In the second course H_2 the acute episode is followed by a period of smoldering disease—chronic active. The third pathophysiological pathway H_3 includes clinical resolution with the maintenance of a chronic carrier state. In fact, there are more than three alternate courses²⁴ but these are not modeled. Elaboration creates a separate class (as discussed on page 80) for each mutually exclusive hypothesis. The result is shown in Figure 3.3.

To those readers who are wondering whether the antibody (anti-HBs) to the hepatitis B surface antigen (HBsAg) should not be made mutually exclusive to the chronic active or chronic persistent courses, I point out that there are cases where the antibody has been detected in these chronic courses but only when the HBsAg antigen has not been detected [43]. THRIPT's knowledge base reflects this by asserting that the end of the period of detectable HBsAg precedes the beginning of the period of detectable anti-HBs.

²³ Again, for demonstration purposes, I have ignored hypotheses which should be considered, even in the absence of additional data. A physician would, in this case, seriously consider the hypotheses of chemical hepatitis.

²⁴ e.g. The chronic active state can, in some cases, progress to recovery or a chronic carrier state and there also is a chronic persistent state.




Figure 3.3: The elaborated hypothesis for hepatitis B

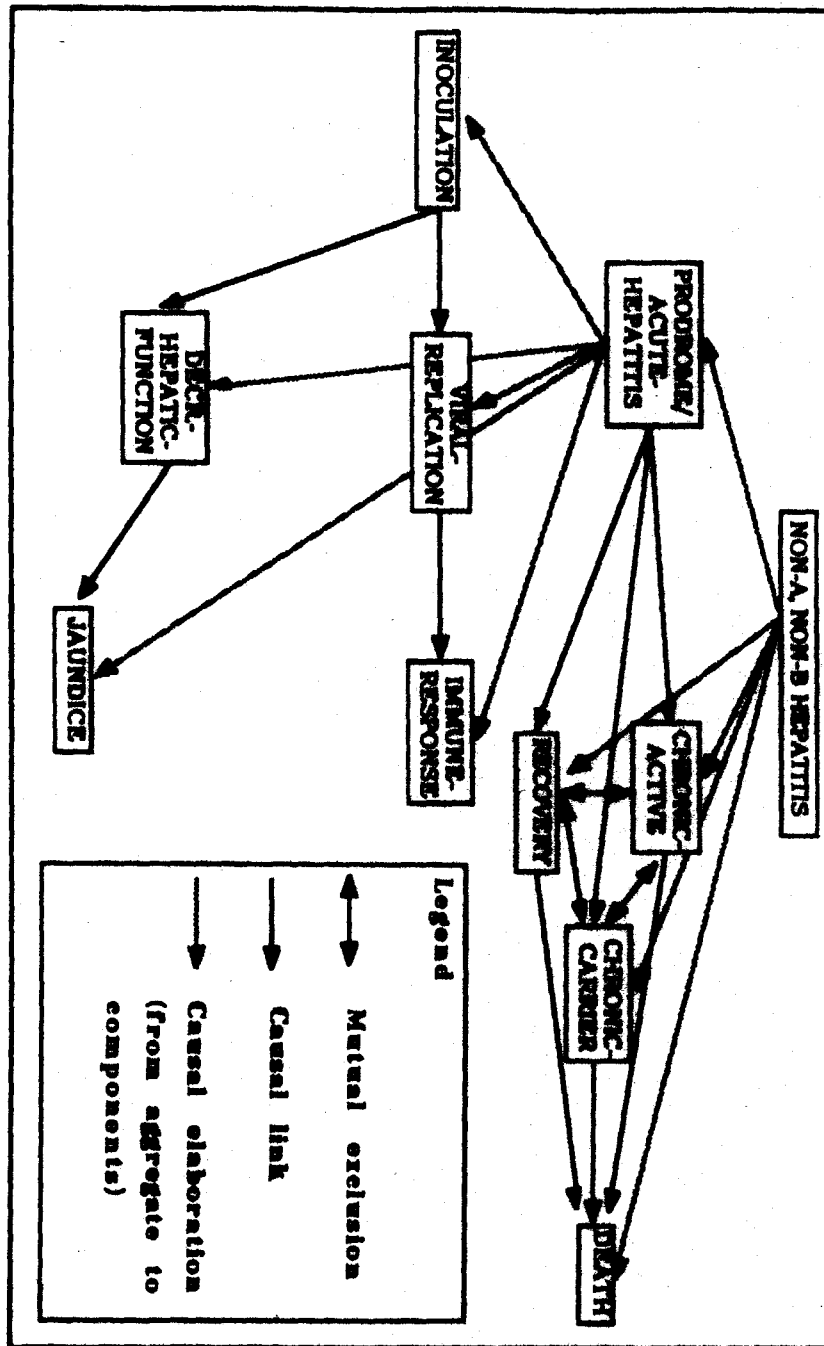


Figure 3.4: Triggered hypothesis for non-A, non-B hepatitis

11. NON-A, NON-B VIRAL HEPATITIS has been elaborated into three mutually exclusive subhypotheses.

The non-A, non-B hepatitis hypothesis has a similar causal structure to hepatitis B (Figure 3.4). At this level of detail, the only difference is that there are no antigens or antibodies. None are modeled as the presumed agent has not been identified.

12. Now beginning instantiation of all subhypotheses...

For each of the mutually exclusive subhypotheses, THRIPTH creates a temporal context. The RRELS are linked to event descriptions²⁵ and therefore RRELS attached to mutually exclusive events are asserted in different contexts. For H_1 , the RRELS associated with the events of the subhypothesis are listed in appendix B.

Each point in each RREL is then given one or more reference set memberships which it shares with all events that share the same causal aggregate(s). THRIPTH then sends the RRELS to TUP to be asserted and their constraints propagated. The result, for H_1 is shown in Figure 3.5. Observe how SOH has taken the "chunking" pattern of the causal aggregation hierarchy of the H_1 hypothesis and translated it into reference sets. Clearly, the prodrome/acute-hepatitis reference set is the largest with five intervals (10 points). The immune response and viral replication reference sets would be of comparable size if I had modeled the many antigens and antibodies that are synthesized during an infection.


Subsequently, temporal assertions of the patient history are added to each context and propagated. Note that since the events of the patient history are anchored with respect to the present, the events of the instantiated hypotheses become anchored as well. For instance, if the patient fits into H_1 population of patients with hepatitis B, the end of jaundice can be expected from two weeks before to 4 weeks after the present (the time when the history was taken).

This process is executed for all subhypotheses.

13. Now beginning hypothesis ranking...

The implementation of the last diagnostic phase is incomplete. Nevertheless, given the instantiated hypotheses in this example, I will take the opportunity to demonstrate that even when atemporal differences between hypotheses can be found, they are often insufficient.

²⁵As implemented, RRELS are added to a hypothesis by using a mouse to "button" an event in a graphical display of the causal graph and then adding a TUP assertion to the temporal assertion slot of that event node in the graph.



The legibility of the individual point and relation labels has been sacrificed to fit the context on one page so that the effect of causal aggregation upon temporal clustering may be observed.

Figure 3.5: Instantiated Hypothesis Without Patient History

To determine whether the present illness was non-A, non-B hepatitis or hepatitis B it would be help THRIPT to find out if and when HBsAg and/or anti-HBs had been detected. To an expert system, without TUP's capabilities at hand, correctly posing the question and then interpreting the answer would be very difficult. Suppose both antigen and antibody had been detected. Whether this represented a bout of hepatitis B prior to the present illness or the resolution of the present illness would depend on the timing with respect to the present of these events. The HBsAg would have to be detected between one to two months after the transfusion and the anti-HBs would have to be detected after the HBsAg.

So, even when a finding is specific to one disease, the temporal dimension may be needed to interpret the finding's absence or presence.

4. Temporal Knowledge Base Engineering

The capability for temporal reasoning significantly improves an expert system's performance: the number of hypotheses considered decreases dramatically (see page 74), diagnostic style becomes more human-like, and greater finesse is achieved in distinguishing between hypotheses. The cost of such an improvement is a much larger effort in knowledge base engineering than required previously. The major part of the effort arises not in the task of constructing the knowledge structures, but in obtaining the required temporal information. In this chapter, I point out some of the difficulties that a knowledge engineer will have in creating temporally-oriented knowledge bases. I also discuss how I have dealt with these obstacles during my experience in using THRIPHT. To illustrate some of the techniques used, I examine (in section 4.4) an excerpt from a medical text and proceed to identify the temporal information it contains.

4.1 Implicit "Common Sense"

It happens, not infrequently, that a neophyte knowledge engineer will set out to represent a body of expertise, expecting that the domain knowledge will be found in some documents used by the domain expert or used to train domain experts. What she soon discovers, especially if the application is broad and realistic, is that a large fraction of the knowledge needed to make the expert system emulate the expert's behavior is not to be found in any document. Nor will it be elaborated by the expert at least during the first few attempts at debriefing, or protocol analysis.

4.1.1 "Common Sense" Temporal Knowledge

One of the reasons for this phenomenon is that there is a lot of "expert's knowledge, that is routinely needed for survival in the world, so commonplace that we are rarely conscious of our use of it. That an object should continue to exist even when it is not perceptible to our senses, that there is a distinction between internal (intrapyschic) events and events external to the individual are not inherently obvious concepts to

an expert system or to a new-born child. To fail to acquire such knowledge of the world is to fundamentally handicap performance. Due to its almost universal use, such expertise is known as "common sense." A large body of temporal knowledge is common sense and is therefore difficult to identify and formalize.

In medicine, in addition to the common sense knowledge that most people share, there is the basic medical knowledge that is acquired early in medical school education and upon which further knowledge acquisition depends. In the professional literature, the basic concepts are usually left implicit and therefore require that the reader be sufficiently trained to understand the explicit information. There is no surprise, then, if an expert system solely built from the information in a medical reference book should fail to simulate a physician's performance. This is also true of the temporal component of medical knowledge.

In a compendium of the infectious diseases, the duration of the incubation period and the duration of the prodrome might be documented with the *implicit* understanding that the onset of the incubation period precedes the onset of the disease prodrome. Similarly, in a discussion of the possible clinical courses of myocardial infarction, the apparently obvious temporal relationships between ischemia and hypoxemia may be omitted.

The pervasiveness and volume of this implicit information means that for every temporal relationship obtained from the domain expert or domain documentation, the knowledge engineer has to consider what additional temporal relationships might be taken for granted and how such information is to be obtained.

4.2 Adherence to TUP Semantics

In chapter 2, I explained how strict adherence to TUP semantics was necessary if TUP's inferences were to be at all meaningful. I discuss here a few of the circumstances in which a knowledge engineer might be led into distorting these semantics.

4.2.1 RREL Bounds

Greater constraint of the temporal context in a medical knowledge base permits more acute differentiation between hypotheses. In the construction of a knowledge base, it is therefore generally desirable to assert RRELS with bounds that are as constrained as possible. It is important however, not to do so in contravention of the RREL bound semantics. For instance, one can find in a medical text book that:

Treatment [of *Haemophilus ducreyi* infection] with tetracyclines or sulfonamides often results in healing in 2 weeks [24, page 215].

The temptation is to construct an RREL as in example 22 with an upper bound of two weeks. Recall however that RREL bounds are the limits on the separation between events within a hypothesis. If the bounds of temporal relations that are "often" observed are asserted in TUP, all consequent temporal inferences no longer have the expected meaning. It then becomes impossible to know what to make of an inconsistency or in fact any temporal inference.

Example 22

```
(RREL ((NAME ANTIBIOTIC-TREATMENT) (TYPE BEGIN-INTERVAL))
      ((NAME HEALING) (TYPE END-INTERVAL))
      (+EPSILON) (+2 WEEKS))
```

Even if probabilistic information could be encoded in RRELS it is not clear how this notion of "often" would be encoded. Instead, the knowledge engineer must determine what the extreme values are for each of the bounds, in this case the maximum delay between onset of treatment and clinical cure. If these extreme values cannot be obtained, and no safe estimate can be made, the knowledge engineer may have to be satisfied with just encoding the precedence information (a lower bound of +EPSILON and an upper bound of +INFINITY). This precedence information alone will go a long way in helping to distinguish hypotheses. Also, as explained in chapter 2, other constraints such as lifespan will cause all events with infinite bounds to be considerably constrained after constraint propagation. As a result, instantiated hypotheses—patient models—do not contain assertions that the patient will heal in infinite time.¹

4.2.2 Loose Bounds and Alternate Contexts

TUP provides two ways of representing uncertain temporal knowledge. One way is to assert a range within which the temporal relation must lie. The other way is to assert several RRELS, one in each context to represent several different possible ranges. Although it is not always clear-cut, there are some reliable guidelines by which one can decide which method to use.

¹This might be accurate only from a rather tasteless perspective.

Setting bounds of a temporal relation is equivalent to asserting that no sub-population of individuals can be identified within the larger population about which more precise (or constrained) assertions could be made. If, however, such distinct subpopulations are identified or at least known to exist, then the bounds particular to each sub-population can be asserted in RRELS in alternate contexts. For example, the delay between the onset of clinical hepatitis and the inoculation with virus, might be some function of the host's immune system activity, nutritional status and numerous other constitutional factors. Until medical science discovers what this function is, or identifies populations of patient with different incubation times, all that can be said of the delay is that it is bounded by 50 to 180 days. The state of knowledge is different for the duration of detectable serum Hepatitis B Surface Antigen (HBsAg). Here, several distinct populations are known to exist and have been identified. In the population of patients whose disease fully resolves, HBsAg disappears approximately within a year of viral exposure. In those unfortunate enough to develop chronic-active hepatitis, the HBsAg levels may persist for several years. It is appropriate in this case to assert an RREL for the duration of HBsAg in each population, each in its respective context.²

Subpopulations of Temporal Percentiles

Contexts permit a limited representation of the *a priori* probabilities of bounds on RRELS. They are of some use in those few occasions when systematic population-wide studies have been made of the temporal progression of events in life and disease. The Denver Developmental Screening Test (DDST) [18] is one such study.

The DDST is designed to provide the pediatrician with a set of landmarks that will permit the evaluation of the developmental progress of the child. The DDST provides, for each behavior, the time by which 25%, 50%, 75% and 90% of the population will have reached the stage at which the behavior can be observed. This information could be represented in THRIPT by having a separate context for each distinct percentile. There would then be a hypothesis for children falling in the first quartile, second quartile, third quartile and so on. If a child was found to begin to say "mama" or "dada" after $8\frac{1}{2}$ months, the percentile knowledge would enable THRIPT to determine that for that particular test of language, the child was in the slowest quartile.

²If this is not clear, consider that once a distinct sub-population is identified, it is a new hypothesis unto itself and therefore requires its own patient model and associated temporal context.

4.3 Volume of Relevant Information

After just a brief experience with temporal knowledge engineering, it becomes clear that for every event/state in the atemporal knowledge base, there are very many temporal relations that have to be encoded. One of the reasons for this is that the global constraint index of a THIRPHT hypothesis has to be quite low to simulate the temporal precision with which human beings distinguish between differing disease chronologies. Usually, individual RREs (before constraint propagation) have relatively high constraint indexes, so that in practice the way to achieve a low global constraint index is to assert several RREs between each event and the other events in the hypotheses so that there is a greater chance of loop creation and consequent constraint propagation.³

Although intuition is often sufficient to determine whether the event positions in a disease chronology have been sufficiently constrained, in practice the knowledge engineer may be misled by common sense knowledge⁴ and overlook some relations that would significantly reduce the global constraint index. I have found in the development of the THIRPHT knowledge bases that it is best to use TUP as an interactive tool to determine if an adequate number of sufficiently constrained temporal relations have been entered. TUP can demonstrate loop creation, measure the global constraint index and verify whether the assertions have produced the expected ordering of events in the hypothesis.

4.4 Example

The example I have chosen to illustrate some of the techniques in temporal knowledge acquisition has been taken verbatim from a widely used textbook of medicine [22]. It is part of the description of the clinical manifestations of pinta—a skin infection with a chronic course. The only modification of the textbook material is the use of bold-face, italic fonts and underlining to identify different types of temporal information. Underlining highlights events, italics denote temporal cues and boldface indicates alternate contexts.

The incubation period in experimental pinta is 7 to 21 days. In man, the initial manifestation is a small papule developing by extension or

³That is, the use of several weak constraints can substitute for fewer strong constraints.

⁴e.g. She may take for granted that virus incubation follows the inoculation.

by coalescence with satellite lesions into a scaly maculopapular lesion. There is regional lymphadenopathy. A generalized erythrosquamous rash develops three to nine months after infection and can be of the "wandering" type... One to three years after the initial lesion, sizable dyschromic macules develop. These late lesions develop from secondary pintides or independently, and pass from slate blue through violet to brown and white—the final achromic phase of the pathogenic process.

This is the clinical course of the *population* of individuals suffering from pinta. This corresponds to THRIPT's composite hypotheses. Since only instantiated hypotheses—patient models—contain events that actually happen to a specific individual, only a patient model can contain temporal relations with respect to the present.⁵ Consequently, in the above excerpt, there are no references to the present.

At first, the natural language equivalents to the events and temporal relations of TUP might be surprising. In TUP's representation, an event is any interval or point in time that is identifiable or distinguished in some way from contiguous temporal points or intervals. In this sense coalescence with satellite lesions is as much an event as the incubation period. Sometimes the references to events are elliptical; in the example, the color slate blue is an ellipsis for "the period during which there are slate blue skin lesions." Care has to be taken to ensure that paraphrases of an event (e.g. late lesions and sizable dyschromic macules) be either translated into the same canonical event in TUP or that the paraphrased events be made temporally equivalent.⁶ The temporal cues are also not those one usually associates with the specification of temporal position; nonetheless, *initial*, *final* as well as the verbs *pass* and *develops* communicate much of the temporal information in the example.

It is apparent that the author of the text quoted above has made some (reasonable) assumptions about the medical knowledge of the reader. The assumptions, however, leave gaps in the temporal knowledge base that the unwary knowledge engineer may miss or have difficulty with. For instance there is in the example the statement that "There is regional lymphadenopathy." To locate this event's temporal position in the knowledge base requires additional (assumed) knowledge—that

⁵Even though this is usually the case, it is not strictly true. A population hypothesis could describe the relation between a clinical syndrome and a specific date (e.g. the high incidence of Parkinsonism after the 1918 outbreak of influenza.) and TUP would compute the position of all events in the hypothesis with respect to the present.

⁶i.e., If the events are intervals, REELs are asserted that make the onset and end of the paraphrases simultaneous.

Lymphadenopathy will follow onset of infection. More than general medical knowledge is needed to further constrain the temporal position of the Lymphadenopathy, but perhaps further constraint is unnecessary for the diagnostic task.⁷

4.5 Finding Reference Sets

We can identify several causal aggregates to represent in a THIRPHT knowledge base and which TUP could then use to drive the generation of reference sets. At the top-most level, there is pinta which is an aggregate of all the possible pathophysiological courses of the disease. A level down, there are the major aggregates of infection, generalised secondary (erythroquamous) rash and the late-lesion phase. We can go yet further and resolve the (acute and initial) infection aggregate into its components: the incubation period, the initial papule, the extension or coalescence with satellite lesions and the scaly maculopapular lesion. A similar decomposition can be made for the other aggregates.

If the domain-dependent organising principle that TUP uses to generate reference sets, reflects a natural decomposition, then the reference sets will have the property of usually containing the more frequently retrieved temporal relations between the closely interrelated events of that domain.

In this example, the more constrained relations tend to occur between events sharing the same reference set. In fact, in the example, events that share reference sets are the only ones with quantitative temporal relations (i.e numerical bounds) and are therefore more constrained than the qualitatively ordered temporal relations of events in separate reference sets. Specifically, the onset and end of incubation share the same reference set: infection; the secondary rash and infection also share the same reference set: Pinta; the late lesions and the infection also share the Pinta reference set.⁸

⁷Perhaps because the duration of regional lymphadenopathy is sufficiently variable and not sufficiently distinct from the duration in other infections to be worthwhile documenting.

⁸Note however that in the example, it is the initial lesion (and not infection) that is explicitly related to the late lesions. I would argue that the author of the pinta article was referring to the lesion as the manifestation of an otherwise clinically silent infection and therefore the two events in this context could be taken to be synonymous. Furthermore, in comparison to the relation between the initial and late lesions which is on the order of years, the initial lesion and infection are contemporaneous. Either synchrony or contemporality would justify my substitution. Meanwhile, as intimated earlier, the Salience Clustering Heuristic is not infallible, and this might be one case where it has slipped

Presumably, the purpose of a medical text is to sufficiently inform the clinician so that he will have at least the minimum information required to make a diagnosis. In this light, the fact that most temporal relations in the example are of the qualitative ordering variety tells us that even relatively unconstrained temporal relations are sufficient for much of the diagnostic task. Only those few relations that are critical in differentiating one disease from other are given numerical bounds. Note that most of the temporal qualifiers in this example only specify the ordering of the onset of intervals, rather than the relative position of interval endings. This may be because the latter ordering provides less diagnostic information.⁹

4.6 Growth in the Number of Contexts

The bold-faced items in the example mark branches into alternate subhypotheses and associated temporal contexts. At every such branch, the number of possible contexts doubles. In the example, there are three branch points and, therefore, eight contexts are required to represent the temporal relations of all the possible pathophysiological progressions.

Large hypotheses, such as I have developed for diseases like hepatitis B, will obviously generate a very large number of contexts. For this reason, if real-time performance is expected of the expert system, then the more likely subhypotheses (and associated temporal contexts) would have to be heuristically selected.

4.7 Summary

The purposes for which this chapter was written include guidance and forewarning in negotiating some of the difficulties of temporally sophisticated knowledge engineering. I end the chapter by emphasizing the forewarning. The example I have used here is one small part of the available knowledge of a relatively uncomplicated disease course. Full-fledged hypotheses, such as are required to diagnose diseases on the order of the hepatitides or ischemic heart disease, are major endeavors requiring at least double the effort of their atemporal versions. This effort arises not in the

up.

⁹Most probably because the duration and therefore ending of the various intervals is much more variable, and therefore less specific, than their onset.

entry of such knowledge but in identifying and acquiring the knowledge. The payoff, of course, is quantitatively and qualitatively improved diagnostic performance.

5. Time In Human Cognition

Well into completing my work on TUP and THRIPT, I became curious as to what was known about the cognitive mechanisms that human beings might employ for temporal reasoning. My interest was quite specific: I wished to see to what extent the fundamental computational limitations of temporal reasoning that I had been faced with in TUP's development (e.g. the explosion of the number of possible temporal deductions) and the solutions I had found, (e.g. reference sets) had their analogies in the human cognitive process. This is not to say that I had a literal-minded expectation of an electrochemical equivalent of the RREL subject to waves of constraint propagation spreading through neuronal networks. I was instead looking for the manifestations of the generic problems associated with temporal reasoning.

There has been a significant effort made in cognitive science in the area of temporal reasoning and representation. Much of it is contested even within the discipline. For instance, in attempting to define the attributes of information that give us our notions of time, several hypotheses have been proposed with varying degrees of success and amounts of supporting experimental evidence. Among these: the amount of memory storage space occupied by an interval of time [44], the effort required to retrieve the event [5] or the intrinsic order of the events represented [39].

5.1 Extent of Temporal Computation

To obtain the temporal position of one event relative to another, *all* event relations that are stored can potentially serve to derive this information. Only a few of these relations will provide the most precise estimate of temporal position or one sufficiently precise to be useful. To find these few useful temporal relations, a large computational investment has to be made. This investment can be made when the temporal data is first accumulated (as in constraint propagation), upon retrieval (as in search) or both (as implemented in TUP). Erickson [15] postulates that in the process of responding to a temporal query, there are distinct subprocesses: retrieving information from memory, coding order information and comparing the retrieved order information with that in the query. Blankenship [3] argues that "encoding information in long-term memory does not automatically incorporate

information about *when* the storage¹ took place." Order judgments may have to be inferred using various context-dependent *sequence rules*. Michon and Jackson [39] point out that there are several cues associated with events that permit the retrieval of order such as causal antecedent/consequent relationships.

Regardless of when and how this computational investment occurs, it grows rapidly with the amount of temporal information (number of events) stored. The critical quantity of information beyond which the computational burden becomes unacceptable depends on the reasoning mechanisms, the representation of temporal information and the performance expected. I feel safe in hypothesizing that this critical quantity is less than the total amount of temporal information stored in the human memory. That is, it would be extremely surprising if the temporal information regarding all events stored in the brain were used to determine the relative position of two particular events.

A survey of the cognitive science literature reveals strong evidence that supports the notion of local temporal reasoning contained within "chunked" groups of events. One line of evidence involves the relatively recent work on cognitive streaming. When the human information processor has to simultaneously handle an overwhelming amount of information, the information is split into several information patterns—cognitive streams. Attention is rapidly switched between streams to maintain an illusion of "real-time" continuity. Bregman [9] in his seminal work on auditory streams provides strong support for the hypothesis that streaming occurs in the auditory process. Among the factors that determine to which of the streams auditory events are added, Bregman describes similarity of spectral composition. Moreover, he finds significant differences in the strength and accuracy of the perception of ordering of auditory events in different streams as compared to events in the same stream. That is, the accuracy of the perception of ordering is more accurate between events that are more closely related.

Michon and Jackson [39] have further investigated cognitive streams and the recall of temporal ordering information as it is affected by event (word) categorization. In their experiments an attempt was made to induce cognitive streaming by presenting the subjects with large numbers of events in a relatively short period. Of particular interest, the experimental evidence seems to demonstrate that the separation of events into cognitive streams does in fact happen, but only if the events are closely related (what Michon and Jackson refer to as *meaning-*

¹Of an event experienced by the subject in real-time. In Blankenship's experiments, the events were jigsaw puzzle tasks.

fulness) which in these experiments, signifies shared word category membership. Significantly, between-category ordering judgments are correct less frequently than within-category ordering judgments. In a similar vein, Blankenship's [3, page 40] experiments show that pair order judgments are much more accurate, and arrived at much faster, if the event pairs have a "context."²

Although I cannot go any further than noting the analogy, it is at the very least intriguing that TUP's solution for managing the temporal relations between large number of events is to group events together in reference sets. Temporal relations between events in different reference sets are not guaranteed to be consistent and furthermore the performance gains expected from reference sets are only accrued if the members of reference sets share some organizing feature or salience—*meaningfulness* or "context."

One of the reasons to be careful in pushing the analogy any further than I have already is that the human performance in temporal reasoning may be just a epiphenomenon of the more general effects of "chunking" observed in the memory storage process.³ For example, Patterson, Maltzer and Mandler [47] have shown that shorter interresponse times (IRTs) occur between items in the same "chunk" than between items in different chunks. The growth of the duration in between-group IRTs depends upon the search mechanism used and is more controversial. Pollio, Richards and Lucas [49] and Patterson *et al.* [47] for instance have developed a probabilistic search model to account for their observed exponential growth of between-group IRTs with output position, whereas McCauley and Kellas [35] observe a linear growth. It is because of these parallels between the performance characteristics of these general retrieval operations and those of temporal recall that it is difficult to isolate the effects of the human temporal reasoning mechanisms.

²In Blankenship's experiments, events with a context were jigsaw puzzle solving tasks in "which the subject saw progressively more and more [of the total puzzle]". The "non-context" puzzle was seen by the subject only one segment at a time.

³Even so, I cannot but point out that TUP's "chunking" of temporal relationships into reference sets is also just a reflection of the more general, atemporal "chunking" of the domain application, be it THERIPHT's causal hierarchy, a KL-ONE classification hierarchy or a planner's goal tree.

5.2 Constraint from Background Temporal Information

When THRIPTH instantiates a hypothesis so that it becomes a patient model, the temporal relations in the patient history may constrain those of the pathophysiological hypothesis and vice versa. This process represents a mapping of certain events in the pathophysiological hypothesis to reported events in the patient history. If the mapping is a correct one, then the temporal information from these two sources may make reciprocal contributions to the knowledge of ordering/position of events from each source. If a patient reports malaise and fever after a blood transfusion (without stating how much earlier the transfusion occurred), if non-A, non-B hepatitis is one of the hypotheses considered, then the period⁴ between malaise and transfusion can be bounded between forty-five and sixty days. Most generally, if TUP is given two different temporal sets of RREL's, and then a few events from one set are *temporally* associated to events in the other set, very often the relations of both sets of RREL's will be constrained if they are asserted in the same context.

A parallel phenomenon may have been observed in human subjects by Guenter and Linton [21]. Subjects were shown sequences of unrelated pictures and then asked to recall the temporal location of these pictures. If a recorded short story that bore no relation to the pictures was provided simultaneously with the projection of the pictures, temporal performance was improved. Guenter and Linton view the story as having provided *temporal tags* for the pictures. My view, which is consistent with Guenter and Linton's interpretation, is that the set of story phrases had more constrained⁵ temporal relationships to one another than the pictures had between each other. Consequently, the story line could constrain the temporal relationships of the pictures, analogous to the manner in which the findings obtained from a patient with hepatitis, constrains the temporal relations of the hepatitis hypothesis.

5.3 Relative Performance of Retrieval Tasks

Of all the temporal retrieval operations that TUP can perform, the one that requires the least computational resources is the determination of the position of one point

⁴Only within that hypothesis.

⁵Since the phrases were related to one another as part of a story line unlike the pictures that were unrelated to one another.

in time relative to another. If the two points are in the same reference set, all that is required is a direct retrieval of the one RREL that links the two points. If the points are in different reference sets, a search is performed. The qualitative⁶ equivalent of the operation is known in the cognitive science literature as an *order* judgment.

The determination of which set of events occur between two time points is performed by the TUP *FINDBETWEEN* function; the cognitive science equivalent is a *lag* judgment. This function is considerably more expensive (for TUP) than range relation retrieval as it involves determining the position of each of the events with respect to the two time points.

An event's ordinal position in a set of events can be obtained using TUP's *FINDPOSITION*—also known as a *position* judgment. It has about one and a half of the cost of a *FINDBETWEEN* computation, as described in section 2.7.2.

This ranking of computational effort is neither inevitable nor necessary. If events were represented by specifying, for each event, a list of those events that occurred before it and those that occurred after it, then position judgment would require the least effort (just counting how many events preceded a particular point in time). Order would be next in computational expense—the point with the shorter “before” list would be the earlier point. Lag would be most expensive (intersection of the “after” list of the earlier point with the “before” list of the later point).

In the development of TUP the choice of temporal relationship for the primary underlying representation, was determined by what I perceived to be the kind of temporal information available (in the medical literature) and by the kind that was most important and most frequently used in making temporal distinctions between (medical) hypotheses.

Experimental evidence points to the similarity in the ordering of difficulty of these tasks in human beings to that in TUP. The work of Jackson and Michon [23], for instance, shows ordering judgments to be less difficult than position and lag judgments. Lag judgments also appear to require less time than position judgments.

Although teleological arguments are not very helpful in supporting hypotheses, the coherence they provide is nonetheless satisfying. In this spirit, one could hypothesize that the mechanisms developed through the biological, evolutionary process would select those cognitive mechanisms that produce the best performance for the kind of temporal reasoning that man most frequently performs. It might be then that recognizing the ordering of event pairs is much more frequently necessary than

⁶Qualitative only that the ordering but not temporal distance is measured.

are lag judgments in going about the business of survival.

5.4 Summary

Artificial intelligence and cognitive science are sister disciplines. Although the methodology and some of the goals may differ, there is a shared interest in the mechanisms of cognition. However it is only in a limited number of areas, such as vision, that the common interest has produced some synergistic interactions with interesting results. My own brief survey of the literature of cognitive science suggests that temporal reasoning is another area that may be ripe for the multidisciplinary approach. In this chapter I touched upon analogies in the performance of TUP and human beings. How this analogy in performance bears on similarities in representation or reasoning mechanisms is not at all apparent, but suggests further investigations to be pursued in this area.

6. Conclusion

6.1 The Problem

Very early in my study of automated diagnostic systems and in particular medical expert systems, it became rather obvious that a crucial component of the diagnostic armamentarium—disease chronology—was not accessible using the available knowledge-engineering tools. As diseases are not static collections of signs, symptoms and pathophysiological states, but dynamic processes with specific temporal patterns, such a deficiency appeared to be a fundamental obstacle to achieving human-like style and human-like performance.

Without access to temporal information about the patient and temporal knowledge of disease, expert systems consider hypotheses that account for the findings but not in the particular order or temporal configuration observed. From the perspective of the human expert such an expert system asks questions that have no apparent bearing on the patient's condition. This reflects that, in the absence of temporal information, many hypotheses are pursued that would be otherwise quickly dismissed by a human expert, or temporally sophisticated expert system. Even if the atemporal discrepancies between the patient data and the expert system's hypotheses eventually become gross enough to permit the exclusion of the incorrect hypotheses, too many questions are asked and too much computational resources are expended in doing so.

Unnecessary generation of questions during an expert system's performance is intrinsically undesirable. Many questions may require diagnostic procedures that have their own cost—discomfort, morbidity, mortality and financial. Also, an expert system that asks too many questions, most of which may seem unrelated to the patient's problem will be unacceptable to the health-care provider because of the attention the system would demand and because such performance would engender a lack of confidence in the diagnostic conclusions. Temporal knowledge by no means eliminates all superfluous questions, but it does prune a large proportion of these.

Lack of temporal representation also implies a lack of a principled distinction between past, present and future. For the diagnostic program, this again leads to the generation of obviously unanswerable questions, for instance regarding the distant future. It also excludes the possibility of a temporally quantified prognosis. For the

medical expert system which plans therapeutic interventions, lack of knowledge of the position of events with respect to the present will lead to absurdities such as making plans to modify the past.

The crux of expert system technology is the explicit representation of all the knowledge that a human expert uses to perform. It is not surprising that in the absence of an explicit representation of an aspect of human reasoning as ubiquitous as temporal reasoning, the expert system should fail to consistently model expert performance. This failure is not restricted to diagnostic and planning performance, but extends to justification, explanation and outcome representation.

6.2 The Plan

In the current literature of knowledge-based systems, there is an almost ritualistic paying of obeisance to the merit of temporal representation or the problems that arise in its absence. In most cases, this is the extent of the concern with the issue. Frequently, an *ad hoc* solution will be described that works for a particular application for most of the test cases. On the other end of the spectrum, a number of coherent and principled temporal logics have been devised that usually provide little guidance or assurance about the use of such systems in implementing the necessary functionality (in AI systems that deal with real-world situations). The design and implementation of TUP and THRIPHT was aimed at providing and demonstrating this functionality without sacrificing generality.

While I developed the tools that would enable expert systems to perform temporal reasoning, three major issues became apparent. The first is the task of supporting the heterogeneity of temporal expression needed for real-world applications. My intent was to do so with a small number of temporal primitives, manipulated by a few temporal operators. Parsimony and simplicity in the underlying representation was emphasized because of the need for a well-understood, uniform method to determine the combination of temporal assertions that produced the most precise estimate of temporal location. The same simplicity would also permit temporal consistency to be readily determined.

The second issue or problem manifested itself in TUP's initial trials. It was apparent that the number of temporal deductions made in realistic applications made temporal reasoning pragmatically infeasible. The obvious solution was to divide the temporal knowledge base into smaller, manageable pieces. The real challenge lay in devising a method for temporal clustering what would minimize

the loss of precision and consistency in the relations retrieved without requiring a large investment of effort on the part of the knowledge engineer.

The last of the three issues and initially the most challenging one, was the integration of temporal reasoning into a medical expert system. It had to be determined if temporal reasoning could be divorced from the other inferences performed in an expert system. One of the themes that recurs in expert systems is that successful system development depends on clear distinctions between the different types of knowledge. Thus, in rule-based systems, the operations of the inference engine are encoded independently¹ and are distinct from the domain knowledge. For the same reasons, it was apparent that temporal reasoning would have to be independent and autonomous from other expert system operations. Such autonomy necessitated the triage of the generic knowledge types in an expert system into two categories: atemporal and temporal.² A skeletal second generation expert system—*TRIAGE*³—had to be developed to investigate the way in which an autonomous temporal reasoner could interact with the atemporal inference mechanisms of such an expert system.

6.3 The Outcome

Once the attempt was made, it was surprising just how clear a distinction there was between the temporal and atemporal elements of the knowledge bases of expert systems. This temporal-atemporal dichotomy was so sharp that it was possible to maintain an independent temporal hypothesis (context) paired with each atemporal domain hypothesis represented. This distinction illuminates some of the previous efforts in knowledge-based systems. For example, in Rieger and Grinberg's [51] causal representation for their "Commonsense Algorithm", several different relations were identified. Many of the relations were further categorised by the temporal relationship, which could be of two types (one shot or continuous). One causal relation was defined as follows:

Action A, or tendency T causes state S, or state change SC to exist, providing gating conditions S1,...,Sn are in effect. For the continuous form, the action's continued presence is required to sustain the state or

¹When it does not, unexpected interactions and type errors eventually are manifested.

²Where the triage was most difficult—recurrent events—the results have been the least successful (see section 6.4).

statechange... For the one shot form, A or T is required only momentarily.

Ignoring the gating conditions³ THRIPT could model a continuous causal link with a vanilla THRIPT causal-link and the following constraints:

```
(RREL ((NAME A) (TYPE BEGIN-INTERVAL))
      ((NAME S) (TYPE BEGIN-INTERVAL))
      (+EPSILON) (+INFINITY))

(RREL ((NAME A) (TYPE END-INTERVAL))
      ((NAME S) (TYPE END-INTERVAL))
      (0 SECONDS) 0 SECONDS))
```

Similarly, a one shot causal link would be modeled with these temporal constraints:⁴

```
(RREL ((NAME A) (TYPE BEGIN-INTERVAL))
      ((NAME S) (TYPE BEGIN-INTERVAL))
      (+EPSILON) (+INFINITY))

(RREL ((NAME A) (TYPE END-INTERVAL))
      ((NAME A) (TYPE BEGIN-INTERVAL))
      (0 SECONDS) 0 SECONDS))
```

The separation of the temporal representation from the causal representation makes it feasible for THRIPT to represent the full (infinite) range of possible temporal configurations between cause and effect instead of restricting these to two categories. Moreover, TUP frees the causal reasoner from the details of managing temporal information by automatically performing those temporal inferences that derive from any combination of cause-effect temporal relationships.

Those aspects of expert system reasoning that were found to be purely temporal and domain-independent were gathered into a package of utilities—TUP. TUP builds its representation on top of two object primitives, the point event and the range relation. Intervals, points, qualitative and quantitative temporal relations, position with respect to the present, common temporal yardsticks, persistence and alternate temporal configurations are all supported by the TUP primitives. The simplicity of the underlying representation makes it easy to determine the precision and consistency of temporal information in the knowledge-base. It also enables TUP to readily

³These could be added to THRIPT's causal links, but in any case their presence is not relevant to the present discussion.

⁴In the second RREL, describing the duration of the action A, the zeroes in the bounds could be replaced by $+\epsilon$ depending on the interpretation one wished to impose.

retrieve the combination of temporal relations that is most precise in calculating the temporal location of an event (with respect to another event).

TUP's clustering scheme—reference sets—was arrived at somewhat serendipitously. Examination of a sample temporal knowledge base revealed that clustering that followed the performance criterion⁵ would parallel the clustering of the atemporal portion of the domain knowledge base. In retrospect, the parallel between the atemporal and temporal decomposition of the knowledge base is not surprising. As explained in chapter 2, both types of clustering are driven by the need to establish correspondence between "salience" or "relevance" and the structure of the knowledge base so as to permit efficient access to the knowledge. THRIPTH illustrates how the atemporal decomposition of the domain knowledge—causal aggregation—can be used to automatically guide the temporal decomposition. It has since become clear that the causal aggregation hierarchy is just one of the knowledge structures that are available to guide temporal clustering. As previously discussed, the structure of the frame-based systems, plans, the hybrid knowledge representation languages, and process descriptions can serve the same purpose.

In addition to increasing the power of an expert system, temporal reasoning considerably loosens the bonds that less expressive expert systems place on the knowledge engineer. Temporal knowledge is so ubiquitous that to attempt to fit it into an *ad hoc* representation of temporal sequence is frustrating. Although TUP's general-purpose temporal representation eliminates this problem, it creates a new obligation for the knowledge engineer—to go out and extract this same ubiquitous temporal information.

6.4 Complications

TUP's design and implementation addresses many of the issues in automated temporal reasoning, with some notable exceptions. Of these the one I feel to be the most important is the lack of an elegant and general utility for representing recurrent events. In the absence of such a utility, I have had to settle for a very limited capability as described in section 2.9.2. Nonetheless, in the same section, I mention some promising methods.

TUP's temporal relation, the *range relation* is devoid of probabilistic information. In many domain applications, systematic temporal probabilistic information is not

⁵i.e. The frequency of retrieval of individual temporal relations.

available. This is fortunate because implementing the capability to reason with probabilistic temporal information is a formidable task. The goals of generality and robustness require, however, that the representation and reasoning mechanisms for such a capability be eventually implemented.

6.5 The Prognosis

Early in the development of expert systems technology, it was recognized that explicit representation of domain-specific knowledge enabled the solution of a large class of problems that were intractable using general purpose problem solvers. From the initial realization that "In the knowledge lies the power" there has developed a corollary axiom (or article of faith): the performance of an expert system is limited and brittle to the extent to which the knowledge representation compels the knowledge engineer to distort the semantics of the domain knowledge. That is, a representation language must mirror the semantics of the type of knowledge it represents.

This is not only true of the expert system's performance, but also of its development and debugging. Consequently there has been a broad effort to create a nosology of the different types of knowledge that an expert uses and therefore that an expert system should explicitly represent. The study of the different types of knowledge in an expert system has included: identifying and representing expert system strategies [11]; explicitly representing the preferences that lead to the goal structure of a decision-maker [63]; explicitly representing the causal aggregation hierarchy of hypotheses [50,46], representing the domain principles that justify the expert performance [54] and explicitly representing spatial relationships [37]. The work that has led to this thesis has been similarly motivated. I have sought to identify that part of the knowledge base that involves temporal information and then gone on to provide an autonomous utility for reasoning about it in a consistent and principled manner. As temporal knowledge is ubiquitous, and particularly so in a patient history, the need for such a capability is great. Nevertheless, like the other types of reasoning just mentioned, temporal reasoning is necessary but not sufficient. My goals will therefore have been well satisfied if the the knowledge gained here becomes part of the knowledge engineer's armamentarium for the development of the next generation of expert systems.

A. Tables and Summaries

A.1 Predicates and Retrieval Functions

The list below summarizes the notation used to describe the predicates and functions. All descriptions omit the context specification, which is an optional argument. All predicates are shown in their relaxed versions. The strict versions include an "S-" prefix to the predicate name.

- $\langle p1 \rangle$, $\langle p2 \rangle$ and $\langle point \rangle$ are point specifications.
- $\langle lb \rangle$, $\langle ub \rangle$ are lower and upper bounds respectively.
- $\langle limit \rangle$ is a single bound.
- $\langle filter \rangle$ is any combination of predicates that returns a boolean value.
- $\langle scope \rangle$ is a list of reference sets.
- $\langle point\ set \rangle$ is a list of points.
- ?X and ?Y are TUP variables.
- $\langle int1 \rangle$ and $\langle int2 \rangle$ are interval specifications.

A.1.1 Predicates

- (BEFORE-P $\langle p1 \rangle$ $\langle p2 \rangle$)
Is $\langle p1 \rangle$ before $\langle p2 \rangle$?
- (AFTER-P $\langle p1 \rangle$ $\langle p2 \rangle$)
Is $\langle p1 \rangle$ after $\langle p2 \rangle$?
- (BEFORE-BY-P $\langle p1 \rangle$ $\langle p2 \rangle$ $\langle lb \rangle$ $\langle ub \rangle$)
Is $\langle p1 \rangle$ before $\langle p2 \rangle$ by the specified amount?
- (AFTER-BY-P $\langle p1 \rangle$ $\langle p2 \rangle$ $\langle lb \rangle$ $\langle ub \rangle$)
Is $\langle p1 \rangle$ after $\langle p2 \rangle$ by the specified amount?
- (WITHIN-P $\langle p1 \rangle$ $\langle p2 \rangle$ $\langle limit \rangle$)
Is $\langle p1 \rangle$ within the $\langle limit \rangle$ distance of $\langle p2 \rangle$?

- (BEFORE-NOW-P <p1>
Is <p1> in the past?
- (AFTER-NOW-P <p1>
Is <p1> in the future?
- (BEFORE-NOW-BY-P <p1> <lb> <ub>)
Is <p1> in the past in the specified range?
- (AFTER-NOW-BY-P <p1> <lb> <ub>)
Is <p1> in the future in the specified range?

A.1.2 Assertions

- (RREL <p1> <p2> <lb> <ub>)
Assert range relation between <p1> and <p2>.
- (ASSERT-INTERVAL <int1>)
Assert the end points of the interval with bounds of 0,+∞.
- (INTREL <int1> <int2> <interval relation>)
Assert the two intervals with RRELS between the end points of the intervals corresponding to the specified interval relation.
- (RelationToPresent <p1> <lb> <ub>)
Assert the RREL between <p1> and the present.

A.1.3 Point Functions

- (GETEVENT <p1> <filter>)
Obtain the point consistent with the <p1> specification if filter is true.
- (FINDBETWEEN <p1> <p2> <scope>)
Obtain the points between <p1> and <p2> that are in scope.
- (FINDPOSITION <point> <point set>)
Obtain the ordinal position of point in point set.

A.1.4 Relation Retrieval

- (RREL <p1> <p2> ?X ?Y)
Return the bounds on the RREL between <p1> and <p2> and also bind the returned values to the TUP variables.

- **(INTREL <int1> <int2> ?X)**
Return a list of interval relations which are consistent with the RRELs between int1 and int2. Also bind the returned list to the TUP variable.
- **(RelationToPresent <p1> ?X ?Y)**
Return the bounds on the RREL between <p1> and the current instance of NOW and bind the returned value to the TUP variable. Every time this query is evaluated a new instance of NOW is generated and its relationship to the current time, obtained from the host computer real-time clock, is asserted. This side-effect is necessary for an accurate answer.
- **(LB-OF <p1> <p2>)**
Return the lower bound on the RREL between <p1> and <p2>.
- **(UB-OF <p1> <p2>)**
Return the upper bound on the RREL between <p1> and <p2>.

A.2 Range Addition

The table below illustrates the rules of range addition for all the different values for the two bounds. l , n and m can be any signed number of seconds. l is distinguished in that it can be any number other than zero.

First bound	Second bound	Range Addition
$+\epsilon$	$-\epsilon$	$-\epsilon$ if the two bounds are lower bounds, $+\epsilon$ if the two bounds are upper bounds.
$+\epsilon$	l	l
0	$+\epsilon$	$+\epsilon$
0	$-\epsilon$	$-\epsilon$
$+\epsilon$	$+\epsilon$	$+\epsilon$
n	m	$n + m$
n	$+\infty$	$+\infty$
n	$-\infty$	$-\infty$
$-\infty$	$-\infty$	$-\infty$
$+\infty$	$+\infty$	$+\infty$
$+\infty$	$-\infty$	$-\infty$ if the two bounds are lower bounds, $+\infty$ if the two bounds are upper bounds.

Table A.1: Range Addition

A.3 Interval to Range Relation Conversion

Interval Relation	Equivalent RRELS
(INTREL {(NAME A)} {(NAME B)} BEFORE)	(RREL ((NAME A) (TYPE END-INTERVAL)) {(NAME B) (TYPE BEGIN-INTERVAL)} {+EPSILON} {+INFINITY})
(INTREL {(NAME A)} {(NAME B)} AFTER)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) {(NAME B) (TYPE END-INTERVAL)} {-INFINITY} {-EPSILON})
(INTREL {(NAME A)} {(NAME B)} DURING)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) {(NAME B) (TYPE BEGIN-INTERVAL)} {-INFINITY} {-EPSILON}) (RREL ((NAME A) (TYPE END-INTERVAL)) {(NAME B) (TYPE END-INTERVAL)} {+EPSILON} {+INFINITY})

Table A.2: Interval Relation To RREL Conversions

Interval Relation	Equivalent RRELs
(INTREL {(NAME A)} {(NAME B)} CONTAINS)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) { (NAME B) (TYPE BEGIN-INTERVAL) } (+EPSILON) (+INFINITY)) (RREL ((NAME A) (TYPE END-INTERVAL)) { (NAME B) (TYPE END-INTERVAL) } (-INFINITY) (-EPSILON))
(INTREL {(NAME A)} {(NAME B)} OVERLAPS)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) { (NAME B) (TYPE BEGIN-INTERVAL) } (+EPSILON) (+INFINITY)) (RREL ((NAME A) (TYPE END-INTERVAL)) { (NAME B) (TYPE BEGIN-INTERVAL) } (-INFINITY) (-EPSILON)) (RREL ((NAME A) (TYPE END-INTERVAL)) { (NAME B) (TYPE END-INTERVAL) } (+EPSILON) (+INFINITY))
(INTREL {(NAME A)} {(NAME B)} OVERLAPPED-BY)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) { (NAME B) (TYPE BEGIN-INTERVAL) } (-INFINITY) (-EPSILON)) (RREL ((NAME A) (TYPE END-INTERVAL)) { (NAME B) (TYPE END-INTERVAL) } (-INFINITY) (-EPSILON)) (RREL ((NAME A) (TYPE BEGIN-INTERVAL)) { (NAME B) (TYPE BEGIN-INTERVAL) } (+EPSILON) (+INFINITY))

Interval Relation	Equivalent RRELS
(INTREL {(NAME A)} {(NAME B)} STARTS)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) {(NAME B) (TYPE BEGIN-INTERVAL)} (0 SECONDS) (0 SECONDS)) (RREL ((NAME A) (TYPE END-INTERVAL)) {(NAME B) (TYPE END-INTERVAL)} (+EPSILON) (+INFINITY))
(INTREL {(NAME A)} {(NAME B)} STARTED-BY)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) {(NAME B) (TYPE BEGIN-INTERVAL)} (0 SECONDS) (0 SECONDS)) (RREL ((NAME A) (TYPE END-INTERVAL)) {(NAME B) (TYPE END-INTERVAL)} (-INFINITY) (-EPSILON))
(INTREL {(NAME A)} {(NAME B)} FINISHES)	(RREL ((NAME A) (TYPE END-INTERVAL)) {(NAME B) (TYPE END-INTERVAL)} (0 SECONDS) (0 SECONDS)) (RREL ((NAME A) (TYPE BEGIN-INTERVAL)) {(NAME B) (TYPE-BEGIN-INTERVAL)} (-INFINITY) (-EPSILON))

Interval Relation	Equivalent RRELS
(INTREL ((NAME A)) ((NAME B)) FINISHED-BY)	(RREL ((NAME A) (TYPE END-INTERVAL)) ((NAME B) (TYPE END-INTERVAL)) (O SECONDS) (O SECONDS)) (RREL ((NAME B) (TYPE BEGIN-INTERVAL)) ((NAME A) (TYPE BEGIN-INTERVAL)) (-INFINITY) (-EPSILON))
(INTREL ((NAME A)) ((NAME B)) EQUALS)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) ((NAME B) (TYPE BEGIN-INTERVAL)) (O SECONDS) (O SECONDS)) (RREL ((NAME B) (TYPE END-INTERVAL)) ((NAME A) (TYPE END-INTERVAL)) (O SECONDS) (O SECONDS))
(INTREL ((NAME A)) ((NAME B)) MEETS)	(RREL ((NAME A) (TYPE END-INTERVAL)) ((NAME B) (TYPE BEGIN-INTERVAL)) (O SECONDS) (O SECONDS))
(INTREL ((NAME A)) ((NAME B)) MET-BY)	(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) ((NAME B) (TYPE END-INTERVAL)) (O SECONDS) (O SECONDS))

A.4 Range Relation to Interval Relation Conversion

RREL Assertion	Consistent INTRELS
(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) ((NAME B) (TYPE BEGIN-INTERVAL)) (+EPSILON) (+INFINITY))	BEFORE OVERLAPS MEETS CONTAINS FINISHED-BY
(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) ((NAME B) (TYPE BEGIN-INTERVAL)) (0 SECONDS) (0 SECONDS))	STARTS STARTED-BY EQUAL
(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) ((NAME B) (TYPE BEGIN-INTERVAL)) (-INFINITY) (-EPSILON))	AFTER OVERLAPPED-BY MEETS FINISHES DURING
(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) ((NAME B) (TYPE END-INTERVAL)) (+EPSILON) (+INFINITY))	BEFORE OVERLAPS MEETS OVERLAPPED-BY FINISHES STARTS FINISHED-BY STARTED-BY CONTAINS CONTAINED-BY

Table A.3: RREL to Interval Relation Conversions

RREL Assertion	Consistent INTRELS
(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) ((NAME B) (TYPE END-INTERVAL)) (O SECONDS) (O SECONDS))	EQUAL MET-BY
(RREL ((NAME A) (TYPE BEGIN-INTERVAL)) ((NAME B) (TYPE END-INTERVAL)) (-INFINITY) (-EPSILON))	AFTER
(RREL ((NAME A) (TYPE END-INTERVAL)) ((NAME B) (TYPE BEGIN-INTERVAL)) (+EPSILON) (+INFINITY))	BEFORE
(RREL ((NAME A) (TYPE END-INTERVAL)) ((NAME B) (TYPE BEGIN-INTERVAL)) (O SECONDS) (O SECONDS))	EQUAL MEETS

RREL Assertion	Consistent INTRELS
<pre>(RREL ((NAME A) (TYPE END-INTERVAL)) ((NAME B) (TYPE BEGIN-INTERVAL)) (-INFINITY) (-EPSILON))</pre>	<p>AFTER OVERLAPS OVERLAPPED-BY CONTAINS DURING MET-BY EQUAL FINISHES FINISHED-BY STARTS STARTED-BY</p>
<pre>(RREL ((NAME A) (TYPE END-INTERVAL)) ((NAME B) (TYPE END-INTERVAL)) (+EPSILON) (+INFINITY))</pre>	<p>BEFORE OVERLAPS MEETS DURING STARTS</p>
<pre>(RREL ((NAME A) (TYPE END-INTERVAL)) ((NAME B) (TYPE END-INTERVAL)) (0 SECONDS) (0 SECONDS))</pre>	<p>EQUAL FINISHES FINISHED-BY</p>
<pre>(RREL ((NAME A) (TYPE END-INTERVAL)) ((NAME B) (TYPE END-INTERVAL)) (-INFINITY) (-EPSILON))</pre>	<p>CONTAINS OVERLAPPED-BY AFTER MET-BY STARTED-BY</p>

B. RRELS of the Annotated Example

B.1 RRELS of a Subhypothesis of Hepatitis B

(RREL ((NAME INOCULATION) (TYPE BEGIN-INTERVAL)) ((NAME INOCULATION) (TYPE END-INTERVAL)) (5 MINUTES) (1 DAY))	(RREL ((NAME RECOVERY) (TYPE BEGIN-INTERVAL)) ((NAME RECOVERY) (TYPE END-INTERVAL)) (+EPSILON) (+INFINITY))
(RREL ((NAME DECR-HEPATIC-FUNCTION) (TYPE BEGIN-INTERVAL)) ((NAME DECR-HEPATIC-FUNCTION) (TYPE END-INTERVAL)) (1 MONTHS) (6 MONTHS))	(RREL ((NAME HEPATITIS B) (TYPE BEGIN-INTERVAL)) ((NAME HEPATITIS B) (TYPE END-INTERVAL)) (+EPSILON) (+INFINITY))
(RREL ((NAME HBSAG) (TYPE BEGIN-INTERVAL)) ((NAME HBSAG) (TYPE END-INTERVAL)) (1 MONTHS) (7 MONTHS))	(RREL ((NAME JAUNDICE) (TYPE BEGIN-INTERVAL)) ((NAME JAUNDICE) (TYPE END-INTERVAL)) (0 SECONDS) (4 WEEKS))
(RREL ((NAME ANTI-HBS) (TYPE BEGIN-INTERVAL)) ((NAME ANTI-HBS) (TYPE END-INTERVAL)) (+EPSILON) (+INFINITY))	(RREL ((NAME HBSAG) (TYPE END-INTERVAL)) ((NAME ANTI-HBS) (TYPE BEGIN-INTERVAL)) (0 SECONDS) (2 WEEKS))
(RREL ((NAME ISSUE-RESPONSE) (TYPE BEGIN-INTERVAL)) ((NAME ISSUE-RESPONSE) (TYPE END-INTERVAL)) (+EPSILON) (+INFINITY))	(RREL ((NAME INOCULATION) (TYPE BEGIN-INTERVAL)) ((NAME JAUNDICE) (TYPE BEGIN-INTERVAL)) (50 DAYS) (6 MONTHS))
(RREL ((NAME VIRAL-REPLICATION) (TYPE BEGIN-INTERVAL)) ((NAME VIRAL-REPLICATION) (TYPE END-INTERVAL)) (1 MONTHS) (6 MONTHS))	(RREL ((NAME DEATH) (TYPE BEGIN-INTERVAL)) ((NAME HEPATITIS B) (TYPE END-INTERVAL)) (0 SECONDS) (0 SECONDS))
	(RREL ((NAME INOCULATION) (TYPE BEGIN-INTERVAL)) ((NAME HEPATITIS B) (TYPE BEGIN-INTERVAL)) (0 SECONDS) (0 SECONDS))

```

(RREL ((NAME DEATH)
      (TYPE BEGIN-INTERVAL))
      ((NAME DEATH)
      (TYPE END-INTERVAL))
      (+INFINITY) (+INFINITY))

(RREL ((NAME HBSAG)
      (TYPE BEGIN-INTERVAL))
      ((NAME INOCULATION)
      (TYPE BEGIN-INTERVAL))
      (-2 MONTHS) (-1 MONTHS))

(RREL ((NAME INOCULATION )
      (TYPE BEGIN-INTERVAL))
      ((NAME VIRAL-REPLICATION)
      (TYPE BEGIN-INTERVAL))
      (1 MONTHS) (2 MONTHS))

(RREL ((NAME VIRAL-REPLICATION)
      (TYPE BEGIN-INTERVAL))
      ((NAME DECR-HEPATIC-FUNCTION)
      (TYPE BEGIN-INTERVAL))
      (4 WEEKS) (8 WEEKS))

(RREL ((NAME DECR-HEPATIC-FUNCTION)
      (TYPE BEGIN-INTERVAL))
      ((NAME JAUNDICE)
      (TYPE BEGIN-INTERVAL))
      (+EPSILON) (+INFINITY))

(RREL ((NAME VIRAL-REPLICATION)
      (TYPE BEGIN-INTERVAL))
      ((NAME IMMUNE-RESPONSE)
      (TYPE BEGIN-INTERVAL))
      (+EPSILON) (+INFINITY))

(RREL ((NAME VIRAL-REPLICATION)
      (TYPE BEGIN-INTERVAL))
      ((NAME HBSAG)
      (TYPE BEGIN-INTERVAL))
      (+EPSILON) (1 DAYS))

(RREL ((NAME IMMUNE-RESPONSE)
      (TYPE BEGIN-INTERVAL))
      ((NAME ANTI-HBS)
      (TYPE BEGIN-INTERVAL))
      (+EPSILON) (+INFINITY))

(RREL
  ((NAME HEPATITIS-B)
  (TYPE BEGIN-INTERVAL))
  ((NAME PRODROME/ACUTE-HEPATITIS)
  (TYPE BEGIN-INTERVAL))
  (50 DAYS) (6 MONTHS))

(RREL
  ((NAME PRODROME/ACUTE-HEPATITIS)
  (TYPE BEGIN-INTERVAL))
  ((NAME PRODROME/ACUTE-HEPATITIS)
  (TYPE END-INTERVAL))
  (3 MONTHS) (6 MONTHS))

```

```

(RREL
  ((NAME PRODROME/ACUTE-HEPATITIS)
  (TYPE END-INTERVAL))
  ((NAME RECOVERY)
  (TYPE BEGIN-INTERVAL))
  (+EPSILON) (+EPSILON))

```

B.2 Notes on the RRELS

1. HBSAG and ANTI-HBS refer to detectable levels of the hepatitis B surface antigen and antibody respectively.
2. IMMUNE-RESPONSE is an aggregate of the causally-connected set of events that are part of the immune system's response to the viral infection. For the purpose of demonstration, I have not represented any of these component events except for the ANTI-HBS in either the RRELS or in the atemporal causal hypothesis.

B.3 Temporal Assertions from the Patient History

(RelationToPresent ((NAME INOCULATION)
(TYPE BEGIN-INTERVAL))
(0 WEEKS) (12 WEEKS))

(RREL ((NAME INOCULATION)
(TYPE BEGIN-INTERVAL))
(NAME JAUNDICE)
(TYPE BEGIN-INTERVAL))
(8 WEEKS) (9 WEEKS))

(RREL ((REFSYSFORM "'45 YEARS OLD'")
(REFSYS AGE)
(TYPE POINT))
(REFSYSFORM "'00:01 A.M., JULY 1ST, 1986'")
(REFSYS CALENDAR)
(TYPE POINT))
(0 SECONDS) (0 SECONDS))

(RREL ((NAME JAUNDICE)
(TYPE BEGIN-INTERVAL))
(REFSYSFORM "'8:00 A.M., JUNE 28TH, 1986'")
(REFSYS CALENDAR)
(TYPE POINT))
(0 SECONDS) (0 SECONDS))

B.4 Notes on the Assertions of the Patient History

1. In the disease hypothesis for hepatitis B, the parenteral introduction of the infectious agent is represented as INOCULATION rather than a blood transfusion.

sion. Although this synonym matching could be done automatically, in the current implementation I have to provide the substitution.

2. Observe, in Figure 3.2, that the CALENDAR mini-expert calculates the temporal distance between the three date points. TUP's constraint propagation algorithm does the rest.

C. Abbreviations

Abbreviation	Definition
RREL	The Range Relation is the TUP primitive for representing the temporal distance between two points.
REFSET	The Reference Set restricts constraint propagation to those points that share set membership.
REFSYS	Reference Systems are commonly used temporal "yardsticks" such as the calendar, age, or stages in human development.
SCH	The Saliency Clustering Heuristic exploits the correspondence between the saliency of information and its accessibility in knowledge structures and the parallel between atemporal and temporal saliency to guide the automatic generation of reference sets.
TUP	The Temporal Utility Package supports assertions and queries to a data base of temporal information.
THRIPHT	Temporal Hypothesis Reasoning In Patient History Taking: a diagnostic medical expert system prototype built to demonstrate the use of TUP.

Table C.1: Abbreviations

References

- [1] James F. Allen.
Maintaining knowledge about temporal intervals.
Communications of the ACM, 26(11):832-843, November 1983.
- [2] G. Octo Barnett.
The application of computer-based medical-record systems in ambulatory practice.
New England Journal of Medicine, 310:1643-1650, 1984.
- [3] Donald Allen Blankenship.
Human Memory for Temporal Sequences.
PhD thesis, University of California, San Diego, 1974.
- [4] H. L. Bleich, R. F. Beckley, G. L. Horowitz, et al.
Clinical computing in a teaching hospital.
N Engl J Med, 312(12):756-64, March 1985.
- [5] Richard A. Block.
Memory and the experience of duration in retrospect.
Memory and Cognition, 2:153-160, 1974.
- [6] Daniel G. Bobrow and Terry Winograd.
An overview of KRL, a knowledge representation language.
In Ronald J. Brachman and Hector J. Levesque, editors, *Readings in Knowledge Representation*, chapter 3, pages 263-285, Morgan Kaufmann Publishers, Inc., Los Altos, California, 1985.
- [7] Ronald J. Brachman, Richard E. Fikes, and Hector J. Levesque.
KRYPTON: a functional approach to knowledge representation.
In Ronald J. Brachman and Hector J. Levesque, editors, *Readings in Knowledge Representation*, chapter 6, pages 411-429, Morgan Kaufmann Publishers, Inc., Los Altos, California, 1985.
- [8] Ronald J. Brachman and James G. Schmolze.
An overview of the KL-ONE knowledge representation system.
Cognitive Science, 9:171-216, 1985.
- [9] Albert S. Bregman.
The formation of auditory streams.
In *Attention and Performance VII*, pages 63-75, Lawrence Erlbaum Associates, 1978.
- [10] B. Chandrasekaran and S. Mittal.

- Conceptual representation of medical knowledge for diagnosis by computer:
MDX and related systems.
In M. Yovits, editor, *Advances in Computers*, pages 217-293, Academic Press,
1983.
- [11] William J. Clancey.
The advantages of abstract control knowledge in expert system design.
In *Proceedings of the National Conference on Artificial Intelligence*,
pages 74-78, 1983.
- [12] R. Davis, B. C. Buchanan, and E. H. Shortliffe.
Production rules as a representation for a knowledge-based consultation
program.
Artificial Intelligence, 8:15-45, 1977.
- [13] Johan de Kleer and John Seely Brown.
A qualitative physics based on confluences.
Artificial Intelligence, 24:7-83, 1984.
- [14] Thomas Dean.
Time Map Maintenance.
Research Report 289, Yale University Department of Computer Science, 1983.
- [15] Leonard Ardel Erickson, Jr.
The Mental Representation of Events.
PhD thesis, Stanford University, November 1974.
- [16] Lawrence Marvin Fagan.
VM: Representing Time-Dependent Relations In A Medical Setting.
PhD thesis, Stanford University, June 1980.
- [17] Kenneth Dale Forbus.
Qualitative Process Theory.
AI-TR 789, Massachusetts Institute of Technology, Artificial Intelligence
Laboratory, 545 Technology Square, Cambridge, MA, 02139, July 1984.
- [18] W. Frankenburg et al.
The newly abbreviated and revised Denver Developmental Screening Test.
Journal of Pediatrics, 99:995, 1981.
- [19] Anne Gardner.
A* - optimal search for an optimal solution.
In Avron Barr and Edward A. Feigenbaum, editors, *The Handbook of
Artificial Intelligence*, chapter C3b, pages 64-66, William Kaufmann, Inc.,
Los Altos, California, 1981.

- [20] Anne Gardner.
Bidirectional search.
In Avron Barr and Edward A. Feigenbaum, editors, *The Handbook of Artificial Intelligence*, chapter C3d, pages 72-73, William Kaufmann, Inc., Los Altos, California, 1981.
- [21] R. Kim Guenther and Marigold Linton.
Mechanisms of temporal coding.
Journal of Experimental Psychology: Human Learning and Memory, 104:182-187, 1975.
- [22] Thorsten Guthe.
Nonsyphilitic treponematoses.
In James B. Wyngaarden and Lloyd H. Smith, Jr., editors, *Cecil Textbook of Medicine*, chapter 14, pages 1584-1589, W.B. Saunders Company, Philadelphia, Pennsylvania, 1982.
- [23] J. L. Jackson, J. A. Michon, and A. Vermeeren.
The processing of temporal information.
In John Gibbon and Lorraine Allan, editors, *Timing and Time Perception*, pages 603-604, The New York Academy of Sciences, 1984.
- [24] Ernest Jawetz, Joseph L. Melnick, and Edward A. Adelberg.
Review of Medical Microbiology.
Lange Medical Publications, Los Altos, California, 15th edition, 1982.
- [25] Kenneth Kahn and G. Anthony Gorry.
Mechanising temporal knowledge.
Artificial Intelligence, 9(1):87-106, 1975.
- [26] E. Yu. Kandrashina.
Representation of temporal knowledge.
In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 346-348, 1983.
- [27] J. P. Kassirer and G. A. Gorry.
Clinical problem solving: a behavioral analysis.
Annals of Internal Medicine, 89:245-255, 1978.
- [28] Isaac S. Kohane.
Temporal Reasoning in Medical Expert Systems.
PhD thesis, Boston University, 1987.
- [29] Benjamin Kuipers.
Commonsense reasoning about causality: Deriving behavior from structure.
Artificial Intelligence, 24:169-204, 1984.

- [30] Benjamin Kuipers.
Getting the envisionment right.
In *Proceedings of the National Conference on Artificial Intelligence*,
pages 209-212, American Association for Artificial Intelligence, August
1982.
- [31] Casimir A. Kulikowski and Sholom M. Weiss.
Representation of expert knowledge for consultation: the CASNET and
EXPERT projects.
In Peter Szolovits, editor, *Artificial Intelligence in Medicine*, Westview Press,
1982.
- [32] William J. Long.
Reasoning about state from causation and time in a medical domain.
In *Proceedings of the National Conference on Artificial Intelligence*,
pages 251-254, 1983.
- [33] William J. Long and Thomas A. Russ.
A control structure for time dependent reasoning.
In *Proceedings of the Eighth International Joint Conference on Artificial
Intelligence*, pages 230-232, 1983.
- [34] J. Malik and T. Binford.
Reasoning in time and space.
In *Proceedings of the Eighth International Joint Conference on Artificial
Intelligence*, pages 343-345, August 1983.
- [35] Charley McCauley and George Kellas.
Temporal aspects of storage and retrieval.
Journal of Experimental Psychology, 102:260-265, 1974.
- [36] Drew McDermott.
Data dependencies on inequalities.
In *Proceedings of the Eighth International Joint Conference on Artificial
Intelligence*, pages 288-289, 1983.
- [37] Drew McDermott.
Finding Objects with Given Spatial Properties.
Research Report 195, Yale University Department of Computer Science, 1981.
- [38] Drew McDermott.
A temporal logic for reasoning about processes and plans.
Cognitive Sciences, 6:101-155, 1982.
- [39] John A. Michon and Janet L. Jackson.

- Attentional effort and cognitive strategies in the processing of temporal information.
In John Gibbon and Lorraine Allan, editors, *Timing and Time Perception*, pages 298-321, The New York Academy of Sciences, 1984.
- [40] Marvin Minsky.
A framework for representing knowledge.
In Patrick Henry Winston, editor, *The Psychology of Computer Vision*, chapter 6, pages 211-277, McGraw-Hill, 1975.
- [41] Sanjay Mittal.
Event-based organization of temporal databases.
In *Fourth National Conference of the Canadian Society for Computational Studies of Intelligence*, pages 1-8, 1982.
- [42] Allen Newell and Herbert A. Simon.
Human Problem Solving.
Prentice-Hall, 1972.
- [43] Robert K. Ockner.
Chronic active hepatitis.
In James B. Wyngaarden and Lloyd H. Smith, Jr., editors, *Cecil Textbook of Medicine*, chapter 7, pages 789-792, W.B. Saunders Company, Philadelphia, Pennsylvania, 1982.
- [44] Robert E. Ornstein.
The 'storage size' metaphor.
In *On the Experience of Time*, chapter 3, pages 37-52, Penguin Books Inc., Baltimore, MD, 1970.
- [45] R. S. Patil, P. Szolovits, and W. B. Schwartz.
Information acquisition in diagnosis.
In *Proceedings of the National Conference on Artificial Intelligence*, pages 345-348, American Association for Artificial Intelligence, 1982.
- [46] Ramesh S. Patil, Peter Szolovits, and William B. Schwartz.
Modeling knowledge of the patient in acid-base and electrolyte disorders.
In Peter Szolovits, editor, *Artificial Intelligence in Medicine*, pages 187-222, Westview Press, Boulder, Colorado, 1982.
- [47] K. E. Patterson, R. H. Meltzer, and G. Mandler.
Inter-response times in categorized free recall.
Journal of Verbal Learning and Verbal Behavior, 10:417-426, 1971.
- [48] Stephen G. Pauker, G. Anthony Gorry, Jerome P. Kassirer, and William B. Schwartz.

- Towards the simulation of clinical cognition: Taking a present illness by computer.
American Journal of Medicine, 60:981-996, 1976.
- [49] H. R. Pollio, S. Richards, and R. Lucas.
Temporal properties of category recall.
Journal of Verbal Learning and Verbal Behavior, 8:529-536, 1969.
- [50] Harry E. Pople, Jr.
Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnostics.
In Peter Szolovits, editor, *Artificial Intelligence in Medicine*, pages 119-190, Westview Press, Boulder, Colorado, 1982.
- [51] Chuck Rieger and Milt Grinberg.
The declarative representation and procedural simulation of causality in physical mechanisms.
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 250-256, 1977.
- [52] Alan Russell, Norris McWirtter, David A. Boehm, et al., editors.
Guinness Book of World Records, page 15.
Sterling Publishing Co, Inc, New York, NY, 1987.
- [53] Elisha P. Sacks.
Qualitative mathematical reasoning.
In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 137-139, 1985.
- [54] William R. Swartout.
XPLAIN: A system for creating and explaining expert consulting programs.
Artificial Intelligence, 21:285-325, 1983.
- [55] P. Szolovits and S. G. Pauker.
Research on a medical consultation system for taking the present illness.
In *Proceedings of the Third Illinois Conference on Medical Information Systems*, pages 299-319, University of Illinois at Chicago Circle, November 1978.
- [56] Peter Szolovits and Stephen G. Pauker.
Categorical and probabilistic reasoning in medical diagnosis.
Artificial Intelligence, 11:115-144, 1978.
- [57] Raúl Valdés-Pérez.
Spatio-temporal reasoning and linear inequalities.

- AIM 875, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, May 1986.
- [58] Steven Vere.
Temporal scope of assertions and window cutoff.
In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 1055-1059, 1985.
- [59] Marc Vilain and Henry Kautz.
Constraint propagation algorithms for temporal reasoning.
In *Proceedings of the National Conference on Artificial Intelligence*, pages 377-382, American Association for Artificial Intelligence, 1986.
- [60] Marc B. Vilain.
The restricted language architecture of a hybrid representation system.
In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 547-551, 1985.
- [61] Marc B. Vilain.
A system for reasoning about time.
In *Proceedings of the National Conference on Artificial Intelligence*, pages 197-201, American Association for Artificial Intelligence, 1982.
- [62] Homer R. Warner.
Computer Assisted Medical Decision-Making.
Academic Press, Inc., New York, New York, 1979.
- [63] Michael Paul Wellman.
Reasoning about preference models.
TR 340, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, May 1985.
- [64] Brian C. Williams.
Doing time: putting qualitative reasoning on firmer ground.
In *Proceedings of the National Conference on Artificial Intelligence*, pages 105-112, American Association for Artificial Intelligence, August 1986.
- [65] T. Winograd.
Procedures as a Representation for Data in a Computer Program for Understanding Natural Language.
AI-TR 235, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 545 Technology Square, Cambridge, MA, 02139, February 1971.

This blank page was inserted to preserve pagination.