

A Temporal Extension to the Parsimonious Covering Theory

Jacques Wainer

*Instituto de Computação, Universidade Estadual de Campinas,
Campinas, SP 13083-970, Brazil
wainer@dcc.unicamp.br
+55 (19) 239-3115 (phone, fax)*

Alexandre de Melo Rezende

*Instituto de Computação, Universidade Estadual de Campinas,
Campinas, SP 13083-970, Brazil
arezende@dcc.unicamp.br
+55 (19) 239-3115 (phone, fax)*

Abstract

In this paper, parsimonious covering theory is extended in such a way that temporal knowledge can be accommodated. In addition to causally associating possible manifestations with disorders, temporal relationships about duration and the time elapsed before a manifestation comes into existence can be represented by a graph.

Precise definitions of the solution of a temporal diagnostic problem, as well as algorithms to compute the solutions are provided. The medical suitability of the extended parsimonious cover theory is studied in the domain of food-borne diseases.

Keywords: Parsimonious cover theory, Temporal abductive diagnosis, Automated medical diagnosis, Temporal reasoning

1 Introduction

Diagnostic reasoning is a complex cognitive process that involves knowledge about a particular domain, general and domain specific heuristics about the diagnostic reasoning itself, and constraints imposed by cognitive limitations of the human diagnosticians. Parsimonious covering theory (PCT) [12] is an attempt to formalize diagnostic reasoning, with the advantage that domain knowledge, domain heuristics, and general diagnostic problem solving methodology are clearly separated from each other.

Briefly, the basic version of PCT defines the domain specific knowledge as a set of disorders (causes), a set of manifestations (effects), and a causal relation between disorders and manifestations. The causal relation associates each disorder with the manifestations it may cause. If one or more of those manifestations are actually present in a diagnostic case, then the disorder may be used to “explain” those manifestations. A particular diagnostic problem is defined by a set of manifestations that are actually observed on the patient, and the solution to that problem is to discover sets of disorders that explain all the manifestations present, with some restrictions on these sets of disorders. A limitation of PCT is that the domain specific knowledge is atemporal, that is, one associates each disorder with a set of manifestations, but it is not possible to specify how these manifestations evolve with time. Because of this atemporality PCT can only be used to solve diagnostic problems in which all relevant symptoms are observable at the moment of diagnosis. In many medical domains, and we expect in other diagnostic domains, that is not the case.

In this paper, PCT is extended in such a way that it is possible to associate with manifestations knowledge about their evolution in time. We call this extension *temporal PCT (t-PCT)*. The need for temporal information in diagnostics systems has been recognized for some time [5,8], and temporal reasoning has been combined with diagnostic reasoning in other works, e. g. [2,3]. This research is the first to combine temporal reasoning with PCT.

The semantics of the causal relation in PCT is that the disorder may cause the manifestation. This means that although the disorder explains the manifestation, that is, if the manifestation is present then the presence of the disorder is a possible explanation for that, the fact that the manifestation may not be present is not taken as an evidence against the disorder.

In a second extension of PCT, we allow for both necessary and possible causal connections between disorders and manifestations. Thus, it is possible to state that a disorder necessarily causes a particular manifestation, or that it just may possibly cause the manifestation. This necessary/possible distinction on the causal relation is called *categorical information*. Other researchers have proposed the inclusion of categorical information in diagnostic reasoning [16,7], but we are particularly interested in the interference between the categorical information and temporal reasoning, which has not been addressed by the works mentioned above. This second extension to PCT that combines categorical and temporal information is called *categorical/temporal PCT (ct-PCT)*.

The next section describes the basic PCT, and it is a summary of [12, Chap. 3]. Section 3 discusses the temporal PCT and section 4 discusses the categorical/temporal PCT. Section 5 reports on the implementation of a diagnostic system for food-borne diseases using ct-PCT, and compares its efficiency with




Fig. 1. Causal network of a diagnostic knowledge base $KB = \langle D, M, C \rangle$.

a standard PCT implementation of the same diagnostic system. Finally section 6 discusses the limitations of the model proposed, and explore some future research topics.

2 Basics of Parsimonious Covering Theory

The basic version of PCT [12] uses two finite sets to define the scope of diagnostic problems (see Figure 1). They are the set D , representing all possible *disorders* d_i that can occur, and the set M , representing all possible *manifestations* m_j that may occur when one or more disorders are present.

The relation C , from D to M , associates each individual disorder to its manifestations. An association $\langle d_i, m_j \rangle$ in C means that d_i may directly cause m_j . Together the set of disorders D , the set of manifestations M , and the causal relation C constitutes the knowledge base KB. More formally, a *knowledge base* is defined as a triple $KB = \langle D, M, C \rangle$.

To complete the problem formulation we need a particular diagnostic *case*. We use M^+ , a subset of M , to denote the set of *observations*, which are manifestations that are present in a particular patient case.

Definition 1 A diagnostic problem P is defined as a pair $\langle KB, Ca \rangle$ where:

- $KB = \langle D, M, C \rangle$ is a knowledge base, with D a finite, non-empty set of elements, called disorders, M a finite, non-empty set of elements, called manifestations, and $C \subseteq D \times M$ is a binary relation called causation.
- $Ca = \langle M^+ \rangle$ denotes the case information, where $M^+ \subseteq M$ is the set of observations.

Disorders (manifestations) are usually denoted by the the letter d (m), possibly supplied with a subscript.

2.1 Solution for a Diagnostic Problem

In order to formally characterize the solution of a diagnostic problem, PCT defines the notion of cover, based on the causal relation C , the criteria for parsimony, and the concept of an explanation (explanatory hypothesis).

Definition 2 For any $d_l \in D$ and $m_j \in M$ in a diagnostic problem P

- $effects(d_l) = \{m_j | \langle d_l, m_j \rangle \in C\}$, the set of manifestation directly caused by d_l ;
- $causes(m_j) = \{d_l | \langle d_l, m_j \rangle \in C\}$, the set of disorders which can directly cause m_j .

The set $effects(d_l)$ represents all manifestations that may be caused by disorder d_l . The set $causes(m_j)$ represents all disorders that may cause manifestation m_j . These functions can be easily generalized to have sets as their arguments.

Definition 3 The set $D_L \subseteq D$ is a cover of $M_J \subseteq M$ if $M_J \subseteq effects(D_L)$.

Definition 4 A set $E \subseteq D$ is an explanation of M^+ for a diagnostic problem iff E covers M^+ , and satisfies a given parsimony criterion.

In the following definition we present some possible parsimony criteria ([17] describes other criteria).

Definition 5

- A cover D_L of M_J is said to be minimum if its cardinality is the smallest among all covers of M_J .
- A cover D_L of M_J is said to be irredundant if none of its proper subsets is also a cover of M_J ; it is redundant otherwise.
- A cover D_L of M_J is said to be relevant if it is a subset of $causes(M_J)$; it is irrelevant otherwise.

In many diagnostic problems, one is interested in knowing all plausible explanations for a case, rather than just a single explanation because they, as alternatives, can affect the course of actions taken by the diagnostician. This leads to the following definition of the problem solution:

Definition 6 The solution of a diagnostic problem $P = \langle KB, Ca \rangle$, designated $Sol(P)$, is the set of all explanations of M^+ .

In this paper we will use irredundancy as the parsimony criterion, as suggested by [12]. If one is interested in developing general algorithms for diagnostic problems, irredundancy seems to be the preferable choice since from the set

of all irredundant explanations one can mechanically generate the set of all minimal explanations (by selecting the sets of minimal cardinality) and the set of all relevant explanations (by systematically adding new disorders to some of the irredundant explanations). It is important to notice that minimal cardinality, which is a natural criterion for parsimony, based on the idea that one should not diagnose more disorders than the necessary, is not a general heuristic, but a domain specific choice. For example, in domains where disorders have different likelihoods or prior probabilities it may be more plausible to say that two fairly common disorders are responsible for a set of observations, than to say that a single extremely rare disorder is the cause.

For example, if the knowledge base is the diagram represented in figure 1, and $M^+ = \{m_2, m_5\}$, then $\{d_4\}$ is a minimum explanation, $\{d_2, d_3\}$ is an irredundant explanation, $\{d_1, d_2, d_4\}$ is a relevant explanation, and $\{\{d_4\}, \{d_2, d_3\}, \{d_1, d_3\}\}$ is the solution for the problem using irredundancy as the parsimony criterion.

2.2 Algorithms for PCT

There are basically two approaches for developing algorithms for PCT, depending on how the set M^+ is presented. The set could be presented *a priori* to the algorithm, in which case it will be said that the algorithm is non-interactive. This seems appropriate in situations when one can monitor all possible manifestations, so that the knowledge of which manifestations are present in the case is readily available. In the second alternative, the observations in M^+ are presented to the algorithm one at a time, possibly as answers to a question posed by the diagnostic system. This approach seems more appropriate in situations where it may be costly to obtain all observations.

Algorithms may also differ in the parsimony criterion used to define an explanation: irredundancy or minimal cardinality. [14] discusses two algorithms that uses minimal cardinality as the parsimony criterion: HT an interactive algorithm, and SOLVE a non-interactive. [12] presents the interactive algorithm BIPARTITE which uses irredundancy as the parsimony criterion, and which will be the basis for the algorithms presented in this paper.

BIPARTITE makes use of generators, a compact representation of alternative explanations for a case. For the sake of completeness, and to be able to use them in our modifications of the BIPARTITE algorithms, we will very briefly describe some concepts and operations on generators. The reader should consult [12] for more details.

If g_1, g_2, \dots, g_m are pairwise disjoint subsets of D , then $G_I = \{g_1, g_2, \dots, g_m\}$ is a *generator*, and the class generated by G_I is $[G_I] = \{\{d_1, d_2, \dots, d_m\} | d_i \in g_i\}$. $G = \{G_1, G_2, \dots, G_N\}$ is a *generator-set* if G_I is a generator, and $[G_I] \cap [G_J] =$

\emptyset .

We define the operations *res*, *div*, *augres*, and *revise*, where G and Q are generator-set, $G_I \in G$ and $Q_J \in Q$ are generators, $S_D \subseteq D$ is a set of disorders, and $q_j \in Q_J$ is also a set of disorders.

Although defined in terms of generator and generator-set, the operation *div* is better understood in terms of sets of explanations. Given a set of explanations for a set of manifestations (M^+), represented as a generator set, and the disorders evoked by a new manifestation m , represented as a set of disorders, the *div* returns the explanations of the original M^+ that would also explain the new manifestation m .

$$\begin{aligned} \text{div}(G, S_D) &= \bigcup_{G_I \in G} \text{div}(G_I, S_D) \\ \text{div}(G_I, S_D) &= \{Q_k \mid Q_k = \{q_{k1}, q_{k2}, \dots, q_{kn}\} \\ &\quad \text{and } q_{kj} = \begin{cases} g_j - S_D & \text{if } j < k \\ g_j \cap S_D & \text{if } j = k \\ g_j & \text{if } j > k \end{cases} \end{aligned}$$

The operation *res* is in some way the dual of *div*, given a set of explanations for m^+ and the disorders evoked by a new manifestation m the *res* returns the explanations of M^+ that did not explain the new manifestation.

$$\begin{aligned} \text{res}(G, S_D) &= \bigcup_{G_I \in G} \text{res}(G_I, S_D) \\ \text{res}(G_I, S_D) &= \begin{cases} \{\{g_1 - S_D, \dots, g_n - S_D\}\} & \text{if } g_i - S_D \neq \emptyset \ 1 \leq i \leq n \\ \emptyset & \text{otherwise} \end{cases} \end{aligned}$$

The operations *div* and *res* are then extended to deal with set of set of disorders (represented as generator and generator-set) as their second argument.

$$\begin{aligned} \text{div}(G, Q_J) &= \bigcup_{G_I \in G} \text{div}(G_I, Q_J) \\ \text{div}(G_I, Q_J) &= \begin{cases} \{G_I\} & \text{if } Q_J = \emptyset \\ \text{div}(\text{div}(G_I, q_j), Q_J - \{q_j\}) & \text{otherwise} \end{cases} \end{aligned}$$

$$\begin{aligned}
res(G, Q) &= \begin{cases} G & \text{if } Q = \emptyset \\ res(res(G, Q_J), Q - \{Q_J\}) & \text{otherwise} \end{cases} \\
res(G, Q_J) &= \bigcup_{G_I \in G} res(G_I, Q_J) \\
res(G_I, Q_J) &= \begin{cases} \emptyset & \text{if } Q_J = \emptyset \\ res(G_I, q_j) \cup res(div(G_I, q_j), Q_J - \{q_j\}) & \text{otherwise} \end{cases}
\end{aligned}$$

The *augres* operation is a modification of the *res* operation so that instead of returning the sets of explanations of M^+ that do not explain the new manifestation m , it adds new disorders to those explanations so that now they also explain $M^+ \cup \{m\}$.

$$\begin{aligned}
augres(G, S_D) &= \bigcup_{G_I \in G} augres(G_I, S_D) \\
augres(G_I, S_D) &= \begin{cases} \{\{g_1 - S_D, \dots, g_n - S_D, A\}\} & \text{if } g_i - S_D \neq \emptyset, A \neq \emptyset \\ \emptyset & \text{otherwise} \end{cases} \\
&\text{where } A = S_D - \bigcup_{i=1}^n g_i
\end{aligned}$$

Thus, given a set of explanations to M^+ , and a set of disorders evoked by a new manifestation m , the set of explanations to $M^+ \cup \{m\}$ can be obtained by a combination of the *div* operation and the *augres* operation. This is done by the *revise* operation:

$$\begin{aligned}
revise(G, S_D) &= Q \cap res(Q', Q) \\
&\text{where } Q = div(G, S_D) \text{ and } Q' = augres(G, S_D)
\end{aligned}$$

2.3 Limitations of PCT

PCT is a conceptually simple and powerful theory of diagnostic reasoning. It clearly separates the role of domain knowledge (sets M , D and principally the relation C), the role of general diagnostic reasoning (the parsimony criteria and the definition of cover), and domain heuristics. This separation allows one to gather and express domain knowledge separately from domain heuristics, as opposed to rule base diagnostic systems [15,18], for example. On the other hand, PCT can also be seen as a limited form of an abductive causal theory [6,7,9].

It has been pointed out that PCT has some limitations to represent more complex forms of causal relationships among disorders and manifestations. The most severe one, for the purpose of this research, is the fact that PCT assumes that two disorders do not interfere with each other. It is not possible to represent that the presence of a disorder will change the manifestations of another disorder, or that if two disorders occur simultaneously they will cause manifestations that none would cause without the presence of the other. An extension of PCT that allows for the representation of the interaction among disorders is discussed in [9].

Another problem of PCT is that the solution of a problem tends to have many alternative explanations. Irredundancy as the parsimony criterion is too weak to significantly reduce the number of alternative explanations, as the experimental results reported in [17] confirm.

An approach for reducing the size of the solution of a diagnostic problem is to add probabilistic information to the causal relation, as shown in [12,13] among others, and compute only the most probable explanations.

Given the lack of probability information about disorders and manifestations in some domains, another useful approach is to develop domain specific heuristics that selects from the set of irredundant explanations a subset of more “plausible” ones. Conceptually, the solution generated by the diagnostic system using irredundancy would be filtered by the domain specific heuristics. Because other explanations can be mechanically generated from irredundant explanations, this approach is at least feasible, although not efficient. Once the appropriate domain specific heuristics have been found, it may be possible to incorporate them directly into the algorithm that generates the solution, improving its efficiency.

3 Temporal PCT

The aim of this paper is to extend PCT so that instead of associating to each disorder a set of manifestation, one could associate an evolution of manifestations. Thus, a knowledge base could state that disorder d_1 causes first m_1 which will last between 2 and 5 days, followed in 7 to 14 days by m_2 which may last an undetermined amount of time, and will be followed at any moment by m_3 , and so on. We accomplish this temporal representation using a graph, in which nodes are manifestations and directed arcs between nodes represent temporal precedence. If there is quantitative information about the duration of the manifestation, it is associated with the corresponding node; if there is quantitative information about the elapsed time between the start of two manifestations, it is associated with the corresponding arc. Furthermore,

quantitative information are not represented as a single number, but as an interval. The knowledge base associates each disorder with one such temporal graph.

A similar representation is used for the case information. It is possible then to state that for a particular patient, manifestation m_4 started sometime between 2 and 3 weeks ago, and it lasted for 1 to 2 days, and that the patient has manifestation m_5 but there is no information on when it started.

3.1 Temporal Representation

Time points, or moments, are the primitive objects to represent temporal information. An *interval* is defined as non-empty convex set of time points, or in other words, as a continuous set of moments. An interval is determined by two time points, its lower and upper extremes, and denoted as $I = [I^-, I^+]$, where I^- is the lower extreme and I^+ the upper extreme. Usually, $I^- \leq I^+$; if $I^- = I^+$ then the interval reduces to a time point, and if $I^- > I^+$ then the interval is empty.

An intervals will be used as a range for temporal measures: to state that a temporal measure must be within an interval $I = [I^-, I^+]$ is to state that the actual value d for that measure must be such that $I^- \leq d \leq I^+$. If the measure is not given precisely, as a single number, but as an interval K , then we will say that the measure interval is compatible with the interval I if there are time points common to both intervals, that is, if $I \cap K \neq \emptyset$. If $I = [I^-, I^+]$ and $J = [J^-, J^+]$, we define the following operations on intervals:

- $I + J = [I^- + J^-, I^+ + J^+]$;
- $I \cap J = [\max(I^-, J^-), \min(I^+, J^+)]$;

and use the abbreviation $I < t$, where t is a time point, to mean $I^+ < t$.

A temporal graph is a *direct, acyclic, transitive*, not necessarily connected graph where the nodes are manifestations. The existence of an arc from m_i to m_j in a temporal graph denotes the fact that the beginning of the occurrence of manifestation m_i must precede the beginning of the occurrence of m_j .

Definition 7 *The temporal graph of a disorder $d_l \in D$, $G_l = (V_l, A_l)$, is a direct, transitive and acyclic graph defined as:*

- $V_l \subseteq M$ is the set of manifestations directly caused by d_l , and
- $A_l = \{(m_i, m_j) \mid \text{the beginning of } m_i \text{ occurs before the beginning of } m_j \text{ when the disorder } d_l \text{ is said to be present}\}$.

The impossibility to define cycles is a major restriction on the expressive power of this temporal representation formalism. In other words, it is not possible to represent recurring events. Nevertheless, this restriction is important since it reduces the complexity of the reasoning process [2].

The *temporal distance* between manifestations and the *duration* of a manifestation are represented by functions on the graph, denoted by DIST and DUR, respectively. The temporal distance function DIST associates an interval $R = [R^-, R^+]$ with each arc of a temporal graph G_l . $\text{DIST}(G_l, (m_i, m_j)) = R$ for $(m_i, m_j) \in A_l$, which we will abbreviate as $\text{DIST}_l((m_i, m_j)) = R$, states that the elapsed time between the beginning of m_j and the beginning of m_i in the temporal graph G_l of d_l must be within the interval R . The duration function DUR associates with each node m_i of a temporal graph G_l an interval J , that specifies that the duration of m_i must be within the interval J .

The transitivity of the temporal graph must be consistently carried over to the DIST function: if $\text{DIST}_l(m_i, m_j) = R_1$ and $\text{DIST}_l(m_j, m_k) = R_2$ then $\text{DIST}_l(m_i, m_k) = R_1 + R_2$.

3.2 Temporal Diagnostic Problem Formulation

Definition 8 *The knowledge base of a temporal diagnostic problem is the tuple $\text{KB} = \langle D, M, G, \text{DIST}, \text{DUR} \rangle$ where D and M are defined as before, G is a set of temporal graphs, each one associated with one disorder of D , DIST and DUR are the temporal information functions defined above.*

In order to represent the case information, we will need a set of observations M^+ , as before, supplemented with two temporal functions BEG^+ and DUR^+ . The function BEG^+ associates an interval with some of the observations in M^+ . $\text{BEG}^+(m) = I$ states that m started at any time within interval I . The origin of the time line for describing BEG^+ is arbitrary, provided the same origin is used in all temporal information for the given case.

Similarly, the function DUR^+ associates a duration interval with some of the observations in M^+ , such that the actual duration of the observation is anywhere within that interval. It is important to notice that the model allows for incomplete knowledge about the observations. Both the beginning and the duration of an observation can be stated as an interval or they may not be stated at all.

Definition 9 *A temporal diagnostic problem P is a pair $\langle \text{KB}, \text{Ca} \rangle$ where KB is defined above, and $\text{Ca} = \langle M^+, \text{BEG}^+, \text{DUR}^+ \rangle$ is the case information.*

One can define the *effects* and *causes* functions in a similar way to def-

inition 2. For example $causes(m) = \{d_l | m \in V_l, \text{ for any temporal graph } G_l = (V_l, A_l) \in G\}$, represents the set of disorders that may cause m .

3.3 Solution for a Temporal Diagnostic Problem

In order to define the solution of a temporal diagnostic problem, a set of concepts about temporal inconsistency needs to be defined. These definitions will allow us to remove from the solution all explanations that contain disorders in which the evolution of manifestations contradicts the evolution of the observations in the case. For example, if for the disorder d_6 the manifestation m_1 precedes m_2 but in the case, the occurrence of m_1 started after the occurrence of m_2 , then one can disregard all explanations that contain d_6 , since it contradicts the temporal information of the case.

Definition 10 For a temporal diagnostic problem P , let $G_l = (V_l, A_l)$, $(m_i, m_j) \in A_l$, and $m_i, m_j \in M^+$. The arc (m_i, m_j) in graph G_l is temporally inconsistent with the case iff

$$(BEG^+(m_i) + DIST_l((m_i, m_j))) \cap BEG^+(m_j) = \emptyset.$$

$(BEG^+(m_i) + DIST_l((m_i, m_j)))$ is according to d_l , the possible range for the beginning of manifestation m_j , or in other words, the set of time points in which disorder d_l expects m_j should begin. On the other hand, $BEG^+(m_j)$ is the range of uncertainty for the real starting point of m_j , that is, all time points in $BEG^+(m_j)$ could have been the moment in which m_j really started. If there is no intersection between these two sets then what the disorder expects are the starting times for the manifestation do not correspond to the real starting times, and the arc (m_i, m_j) should be considered temporally inconsistent with the case.

Definition 11 For a temporal diagnostic problem P , let $G_l = (V_l, A_l)$ be the temporal graph of a disorder $d_l \in D$. The disorder d_l is temporally inconsistent with the case $Ca = \langle M^+, BEG^+, DUR^+ \rangle$ iff

- there exists at least one arc $(m_i, m_j) \in A_l$ temporally inconsistent with respect to the case, or
- there exists at least a node $m \in V_l$, such that, $m \in M^+$ and $DUR_l(m) \cap DUR^+(m) = \emptyset$.

Thus, a disorder is temporally inconsistent with the case information, if it has a temporally inconsistent arc, or if the range for the duration of one of its manifestations does not agree with the range for the duration of the corresponding observation.

Finally, based on the above definitions, we formalize the notions of temporally consistent explanation and temporally consistent solution.

Definition 12 *A set $E \subseteq D$ is said to be a temporally consistent explanation of the case for a temporal diagnostic problem P iff*

- (i) E covers M^+ , and
- (ii) E satisfies a given parsimony criterion, and
- (iii) for all $d_i \in E$, d_i is not temporally inconsistent with the case.

Definition 13 *The temporally consistent solution of a temporal diagnostic problem $P = \langle \text{KB}, \text{Ca} \rangle$, designated by $\text{Sol}(P)$, is the set of all temporally consistent explanations of the case.*

3.4 Algorithm

We present here an interactive algorithm that computes all explanations to a temporal diagnostic problem. The algorithm is a modification of the BIPARTITE algorithm in [12]. The important aspect of the algorithm is that temporal consistency is not implemented as a filter, that is, it is not applied after the original BIPARTITE algorithm has generated the solution, but it is incorporated very early into the process of merging the causes on the “new” observation into the set of current explanations. Thus the algorithm has to deal with smaller sets of explanations.

function t-BIPARTITE(KB)

variables

m : **manifestation**; (* new observation *)
 $hypothesis$: **generator-set**; (* all explanations *)
 D_C , (* consistent disorders *)
 D_I , (* inconsistent disorders *)
 H , (* disorders evoked by m *)
 I_B , (* inconsistent disorders due to BEG *)
 I_D : **set-of-disorder**; (* inconsistent disorders due to DUR *)
 M^+ : **set-of-manifestation**;
 BEG^+ ,
 DUR^+ : **function**;

```

1  begin
2     $hypothesis = \{\emptyset\}$ ;
3     $D_C = \emptyset$ ;
4     $D_I = \emptyset$ ;
5     $M^+ = \emptyset$ ;
6    while MoreObservations do
7       $I_B = \emptyset$ ;
```

```

8       $I_D = \emptyset$ ;
9       $m = \text{NextObservation}$ ; (* obtain next observation *)
10      $M^+ = M^+ \cup \{m\}$ ;
11      $H = \text{causes}(m) - D_I$ ;
12     if  $\text{DUR}^+(m)$  is defined
13       then
14          $I_D = \{d_l | d_l \in H, \text{ and } \text{DUR}_l(m) \cap \text{DUR}^+(m) = \emptyset\}$ ;
15     endif
16     if  $\text{BEG}^+(m)$  is defined
17       then
18          $I_B = \text{temp\_inc}((H - I_D) \cap D_C, M^+, m)$ ;
19     endif
20      $\text{hypothesis} = \text{res}(\text{hypothesis}, I_B \cup I_D)$ ;
21      $D_I = D_I \cup I_B \cup I_D$ ;
22      $D_C = (D_C \cup H) - (I_B \cup I_D)$ ;
23     if  $(H - D_I) = \emptyset$  or  $(\text{hypothesis} = \emptyset \text{ and } M^+ \neq \emptyset)$ 
24       then
25         return nil (* there is no consistent explanation *)
26     else
27        $\text{hypothesis} = \text{revise}(\text{hypothesis}, H - D_I)$ ;
28     endif
29   endwhile
30   return hypothesis
31 end.

```

The function *temp_inc* in line 18 is defined as follows: given a set of disorders S_D and a new manifestation, the function returns the disorders from the set S_D that are temporally inconsistent with the new manifestation. In order not to clutter the notation, we assume that the knowledge base $\text{KB} = \langle D, M, G, \text{DIST}, \text{DUR} \rangle$ and the case information $\text{Ca} = \langle M^+, \text{BEG}^+, \text{DUR}^+ \rangle$ are globally accessible to this function.

$$\begin{aligned}
\text{temp_inc}(S_D, m) = & \\
& \{d_l \in S_D \mid \text{there exists } (m_i, m) \in A_l \text{ and } m_i \in M^+ \text{ and} \\
& \quad (\text{BEG}^+(m_i) + \text{DIST}_l((m_i, m))) \cap \text{BEG}^+(m) = \emptyset \\
& \quad \text{or} \\
& \quad \text{there exists } (m, m_j) \in A_l \text{ and } m_j \in M^+ \text{ and} \\
& \quad (\text{BEG}^+(m) + \text{DIST}_l((m, m_j))) \cap \text{BEG}^+(m_j) = \emptyset \\
& \quad \}
\end{aligned}$$

The functions `MoreObservations` and `NextObservation` are entry-points for the module that interacts with the patient, probably through the physician, asking questions about the presence of manifestations. In order to ask effective

questions this module must have access to current set of explanations, the knowledge base and very likely will use domain specific heuristics to select the next question to ask.

At the beginning of a new cycle, a new observation (m) is entered (line 9), and the disorders known to be inconsistent with the temporal information of the case so far (D_I) are removed from the disorders evoked by m (line 11), resulting in a set of consistent disorders evoked by the new observation (H). If there is duration information for the new observation, the set of disorders in H that are inconsistent with the duration information is collected in I_B (line 14); if there is starting time information for the new observation, the set of disorders in H which are inconsistent with this information are collected in I_B (line 18). Explanations that contain such disorders are eliminated from the current hypothesis (line 20). The sets of all inconsistent and consistent disorders so far are updated (lines 21 and 22), and the set of explanations is updated to include the consistent disorders evoked by the new observation (line 27).

3.5 Discussion

We have presented our first extension to PCT, which adds temporal representation of manifestations and observations to the original PCT. The temporal representation used here is similar to the ones used by other researchers both in medical domains [2,3] and robotics [4]. The temporal inconsistency criterion is equivalent to the one described in [3].

As discussed earlier, this temporal representation allows for many forms of uncertainty. Time information may be expressed as intervals or may not be expressed at all, both for the knowledge base and for the case. In fact, the t-PCT is a true extension of the original PCT, since by not providing any temporal information one has both a PCT knowledge base and a PCT case, and in this case the definition of a solution for a temporal diagnostic problem will coincide with the PCT's definition of solution for a diagnostic problem.

This true extension property is a positive trait since many diagnostic domains (including some medical domains) are atemporal in the sense described above, and t-PCT could be the appropriate diagnostic theory for them as well. But the true extension property restricts some possible useful forms of uncertainty in describing the case. For example, it is not possible to state that an observation has already occurred and that it is not present anymore, without explicitly stating its starting time and duration.

4 Categorical/Temporal PCT

The semantics of the causal relation in PCT (and t-PCT) is that the disorder may cause the manifestation. An important extension is to distinguish between a possible causation and a necessary causation [7]. For example, botulism may cause nausea and vomiting, but necessarily causes some form of paralysis. This distinction between necessary causation and possible causation will be called *categorical information*.

By taking into consideration this categorical information, it is possible to reduce the number of explanations for a particular diagnostic problem: if d_1 necessarily causes m_2 , and if m_2 is not one of the manifestations observed in the diagnostic case, then no explanation for the case can contain the disease d_1 . We call this reasoning, *categorical rejection*. When a temporal dimension is added, one has to be careful about categorically rejecting a disorder. Even if d_1 necessarily causes m_2 , and if m_2 is not observed in the case, one should not categorically reject d_1 without checking whether m_2 had time to develop, given the current stage of the disorder d_1 .

4.1 Problem Formulation and its Solutions

In order to provide categorical information, the function POSS is added to the knowledge base. It attributes to each node of each temporal graph either the label N , for necessary, or the label P , for possible; $\text{POSS}(G_l, m_j) = N$, abbreviated as $\text{POSS}_l(m_j) = N$, states that disorder d_l necessarily causes the manifestation m_j .

In categorical diagnostic problems, one is interested in manifestations known to be absent in the diagnostic case, called *negative observations*. M^- , the set of negative observations, and I_{now} , the time point that represents the moment of diagnosis, are added to M^+ , BEG^+ , DUR^+ as the components of the patient case information, Ca.

Definition 14 Let $P = \langle \text{KB}, \text{Ca} \rangle$ be a categorical diagnostic problem and $G_l = (V_l, A_l) \in G$. The disorder d_l is categorically inconsistent with the case iff there exists $m \in V_l$ such that $\text{POSS}_l(m) = N$, and $m \in M^-$ and

- (i) there exists $(m, m_j) \in A_l$, such that $m_j \in M^+$, or
- (ii) there exists $(m_i, m) \in A_l$, such that, $m_i \in M^+$ and $\text{BEG}^+(m_i) + \text{DIST}(m_i, m) \leq I_{\text{now}}$.

For a disorder d_l to be categorically inconsistent a necessary manifestation must not be present in the case ($\text{POSS}_l(m) = N$ and $m \in M^-$), but also

there must have been enough time for the manifestation to occur. In the first case above, a later manifestation has already occurred ((m, m_j) in A_l and $m_j \in M^+$), so we can be sure that m should already have occurred. In the second case above, that is warranted because there has been enough time ($\text{BEG}^+(m_i) + \text{DIST}(m_i, m) \leq I_{\text{now}}$) since the occurrence of a preceding manifestation that did occur in the case ($(m_i, m) \in A_l$ and $m_i \in M^+$) and the necessary manifestation that is not present.

Finally, we define an explanation of a categorical diagnostic problem.

Definition 15 *A set $E \subseteq D$ is said to be a consistent explanation of the case for an categorical diagnostic problem $P = \langle \text{KB}, \text{Ca} \rangle$ iff*

- E covers M^+ , and
- E satisfies a given parsimony criterion, and
- for all $d_l \in E$, d_l is not temporally inconsistent, and
- for all $d_l \in E$, d_l is not categorically inconsistent.

4.2 Algorithm

We present below an algorithm that interactively solves a categorical/temporal diagnostic problem. The algorithm makes use of two auxiliary data structures CCR_1 and CCR_2 , which mainly store disorders that are *candidates for categorical rejection*, that is, disorders that have a necessary manifestation not present in the case, but cannot be categorically rejected because the conditions (i) or (ii) in definition 14 are not true yet. CCR_1 stores disorders for which the condition (i) in definition 14 was not verified and CCR_2 stores disorders for which the other condition was not verified.

CCR_1 is a set of elements of the form $\langle d, \{m_1, \dots, m_k\} \rangle$ where d is a candidate for categorical rejection and $\{m_1, \dots, m_k\}$ are disorders that should occur after a necessary manifestation of d that is not present in the case. CCR_2 is a set of elements of the form $\langle d, \{(m_{b1}, m_{n1}) \dots (m_{bk}, m_{nk})\} \rangle$, where each (m_{bi}, m_{ni}) is an arc of the temporal graph for d , and each m_{ni} is a necessary manifestation of d that is not present in the case, m_{bi} is a manifestation for which the system has no knowledge whether it occurred in the case or not, and that should have occurred before the corresponding m_{ni} .

The operations on CCR_1 and CCR_2 are defined below. *cat_rej* returns the set of disorders that can be categorically rejected once the manifestation m is known to have occurred. *remove* removes a set of disorders from the corresponding list, and *add* adds a new set of disorders to the list. For these operations we assume that the knowledge base and the case information are

globally accessible.

$$cat_rej(CCR_1, m) = \{d_l | \langle d_l, M_l \rangle \in CCR_1 \text{ and } m \in M_l\}$$

$$cat_rej(CCR_2, m, t_{now}) = \{d_l | \langle d_l, A_l \rangle \in CCR_2 \text{ and there exists} \\ (m, m_{ni}) \in A_l \text{ and } BEG^+(m) + DIST_l((m, m_{ni})) \leq t_{now}\}$$

$$remove(CCR, S_D) = \{\langle d, X \rangle | \langle d, X \rangle \in CCR \text{ and } d \notin S_D\}$$

$$add(CCR, S) = \{\langle d, X \rangle | \langle d, X \rangle \in CCR \text{ and } \langle d, Y \rangle \notin S\} \cup \\ \{\langle d, Y \rangle | \langle d, Y \rangle \in S \text{ and } \langle d, Y \rangle \notin CCR\} \cup \\ \{\langle d, X \cup Y \rangle | \langle d, X \rangle \in CCR \text{ and } \langle d, Y \rangle \in S\}$$

function ct-BIPARTITE(KB)

variables

m: **manifestation**; (* new observation *)
hypothesis: **generator-set**; (* all explanations *)
D_C, (* consistent disorders (temp. and categ.) *)
D_I, (* inconsistent disorders (temp. and categ.)*
H, *H₁* (* disorders evoked by *m**)
I_B, (* inconsistent disorders due to BEG *)
I_D, (* inconsistent disorders due to DUR *)
I_C : **set-of-disorder**; (* categorically rejected disorders *)
CCR₁, *CCR₂* :**set**;
M⁺: **set-of-manifestation**; (* observations *)
BEG⁺, *DUR⁺*: **function**;
t_{now}: **time point**; (* now *)
BEG⁺,
DUR⁺: **function**;

1 **begin**

2 *hypothesis* = { \emptyset };

3 *D_C* = \emptyset ;

4 *D_I* = \emptyset ;

5 *CCR₁* = \emptyset ;

6 *CCR₂* = \emptyset ;

7 *M⁺* = \emptyset ;

8 *t_{now}* = Now;

9 **while** MoreObservations **do**

10 *m* = NextObservation;

11 *H* = *causes*(*m*) - *D_I*;

12 **if** NextObservation.status = present (* *m_j* ∈ *M⁺* *)

13 **then**

14 *M⁺* = *M⁺* ∪ {*m*}

```

15      $I_B = \emptyset;$ 
16      $I_D = \emptyset;$ 
17      $I_C = cat\_rej(CCR_1, m) \cup cat\_rej(CCR_2, m, t_{now});$ 
18      $CCR_1 = remove(CCR_1, I_C);$ 
19      $CCR_2 = remove(CCR_2, I_C);$ 
20     if  $DUR^+(m_j)$  is defined
21         then
22              $I_D = \{d_l | d_l \in (H - I_C), \text{ and } DUR_l(m) \cap DUR^+(m) = \emptyset\};$ 
23         endif
24     if  $BEG^+(m_j)$  is defined
25         then
26              $I_B = temp\_inc((H - (I_D \cup I_C)) \cap D_C, M^+, m);$ 
27         endif
28          $hypothesis = res(hypothesis, I_C \cup I_B \cup I_D);$ 
29          $D_I = D_I \cup I_C \cup I_B \cup I_D;$ 
30          $D_C = (D_C \cup H) - (I_C \cup I_B \cup I_D);$ 
31         if  $(H - D_I) = \emptyset$  or  $(hypothesis = \emptyset$  and  $M^+ \neq \emptyset)$ 
32             then
33                 return nil (* there is no consistent explanation *)
34             else
35                  $hypothesis = revise(hypothesis, H - D_I);$ 
36             endif
37     else (* $m_j \in M^{-*}$ *)
38          $H_1 = \{d_l | d_l \in H, \text{ and } POSS_l(m) = N\};$ 
39          $I_C = \{d_l | d_l \in H_1 \text{ and there exists } (m, m_k) \in A_l \text{ such that } m_k \in M^+\}$ 
40              $\cup \{d_l | d_l \in H_1 \text{ and there exists } (m_i, m) \in A \text{ such that}$ 
41                  $BEG^+(m_i) + DIST_l((m_i, m)) \leq t_{now}\};$ 
42          $hypothesis = res(hypothesis, I_C);$ 
43          $D_I = D_I \cup I_C;$ 
44          $D_C = D_C - D_I;$ 
45          $CCR_1 = add(CCR_1, \{\langle d_l, S_M \rangle | d_l \in H_1 - I_C \text{ and}$ 
46              $S_M = \{m_k | (m, m_k) \in A_l\}\});$ 
47          $CCR_2 = add(CCR_2, \{\langle d_l, S_A \rangle | d_l \in H_1 - I_C \text{ and}$ 
48              $S_A = \{(m_i, m) | (m_i, m) \in A_l\}\});$ 
49         if  $hypothesis = \emptyset$  and  $M^+ \neq \emptyset$ 
50             return nil
51         endif
52     endif
53 endwhile
54 return  $hypothesis$ 
55 end.

```

In its main loop, the algorithm is divided into two segments: lines 14 to 36 process a new observation ($m \in M^+$), while lines 38 to 51 process a negative observation ($m \notin M^+$). If the manifestation is present then line 17 collects into

the set I_C all disorders in CCR_1 and CCR_2 that indeed became categorically inconsistent by the presence of m . Lines 18 and 19 update the lists CCR_1 and CCR_2 by removing the categorically rejected disorders from them. Lines 20 to 36 basically repeat the correspondent segment of code in the algorithm t-BIPARTITE, taking also into consideration the categorically inconsistent disorders.

When the manifestation m is not present ($m \in M^-$), the algorithm collects in H_1 the set of disorders evoked by m , which has m as a necessary manifestation. From this set, the categorically inconsistent disorders are collected in I_C , which is used to update the current hypothesis (line 42), update the set of inconsistent disorders (line 43), and update the set of consistent disorders (line 44). Finally the sets CCR_1 and CCR_2 are updated with the disorders in H_1 which are not yet categorically inconsistent with the case.

5 Diagnosis of food-borne disease using ct-PCT

Food-borne diseases [10, Chap. 81] result from the ingestion of food contaminated with pathogenic microorganisms, toxins or chemicals, and their symptoms are primarily gastrointestinal or neurological. This domain is well suited to test both t-PCT and ct-PCT because not only temporal information about the evolution of symptoms and categorical information are available, but they are necessary for a correct diagnosis. For example: *Staphylococcus aureus* and short-incubation *Bacillus cereus* are the only possible bacterial causes for nausea and vomiting occurring within 1 to 6 hours after the ingestion of the contaminated food. On the other hand, a patient with botulism (*Clostridium botulinum*) will only have nausea and vomiting in 18 to 36 hours after the ingestion. Not taking into consideration the elapsed time from ingestion to the symptoms would result in incorrect diagnosis.

Furthermore, food-borne diseases also shows examples in which the categorical rejection must be performed carefully because of temporal considerations. For example the ingestion of poisonous mushrooms from the species *Amanita phalloides*, *A. virosa*, and *A. verna*, will necessarily cause in 6 to 24 hours abdominal cramps and diarrhea, which will last for up to 24 hours, followed by a 1 to 2 days period of no symptoms, followed by hepatic and renal failures (which in almost 50% of the cases leads to death). One should not categorically reject this disease for a patient that is not showing signs of renal failure without taking into consideration whether there has been enough time for that symptom to develop.

We implemented a diagnostic system for food-borne diseases in order to:

- test the theory.
- verify the adequacy of the temporal representation for representing diseases.
- compare the precision and efficiency of PCT and ct-PCT (and the corresponding algorithms) in solving some diagnostic problems in this domain.

In particular it was not our goal to construct a diagnostic system to be used in clinical settings. The implemented system has not been verified and validated appropriately for clinical use.

5.1 The Knowledge Base

The knowledge base contained 28 of the food-borne infections described in [10, Chap. 81], amounting to around 60 different symptoms. The whole knowledge base was developed in four days, based mainly on that medical manual. A specialist was consulted once during the development phase, mainly to provide categorical information on the symptoms of each disease, since such information was not always available (or it was unclear) in the manual. The specialist also verified the temporal graphs for some of the diseases. The total time of consultation with the specialist was around two hours.

Food-borne diseases usually have a simple temporal structure: the event of ingestion of the contaminated food and a set of cotemporal symptoms that occur after the incubation period. Figure 2 represents the temporal graphs of the symptoms caused by *Y. enterocolitica* and by the *A. phalloides*, *A. virosa*, and *A. verna* mushrooms.

The *Y. enterocolitica* graph is typical of food-borne diseases. The manifestation m_0 is not really a manifestation but the event of ingestion. The necessary manifestations are fever (m_1), abdominal cramps (m_2), and mesenteric adenitis (m_3); and the possible manifestations are nausea (m_4) and vomiting (m_5). The manual does not provide information about the elapsed time between ingestion and manifestations m_3 , m_4 , and m_5 , so we define the DIST function for those arcs to be $[0, \infty]$, which only means that those manifestations occur any time after m_0 . The duration of all manifestations is [1d,28d].

The graph for *A. phalloides* is more complex from the temporal point of view. The arcs from m_0 to m_8 and from m_0 to m_9 are not represented in the graph. m_0 again is the ingestion event, m_6 and m_7 are abdominal cramps and diarrhea, and m_8 and m_9 are renal and hepatic failures. The durations of m_6 and m_7 are [6h,24h] and the duration of m_8 and m_9 are not defined, and thus set to be $[0, \infty]$. Other diseases, like *paralytic shellfish poisoning (PSP)* have a temporal graph similar to the one for *Y. enterocolitica*, but with 10 necessary

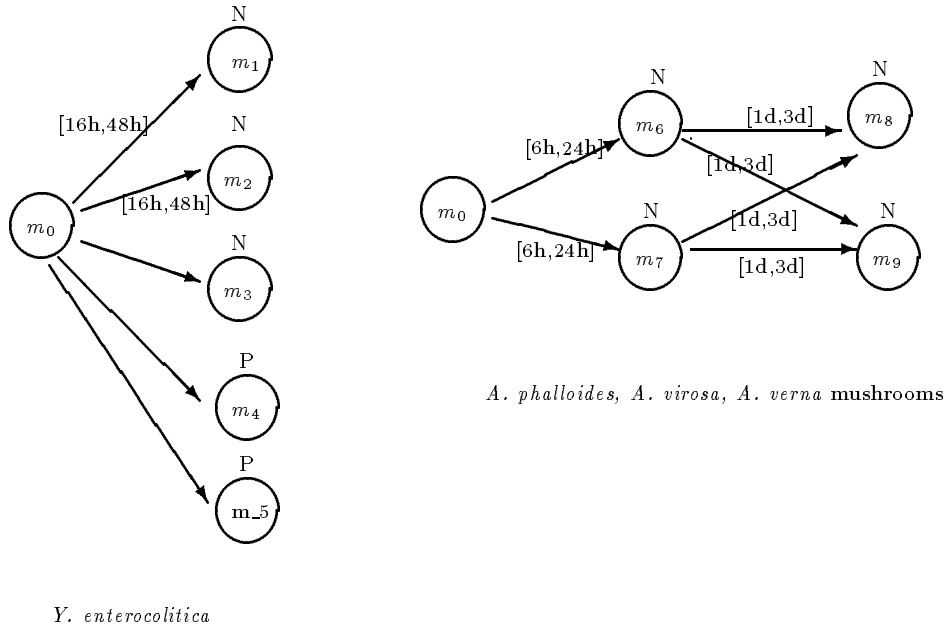


Fig. 2. Temporal graphs of some food-borne diseases.

symptoms and 25 possible symptoms.

We encountered two examples of diseases for which ct-PCT was inadequate to represent the development of the manifestations or the causal relation between the disease and its manifestations. This two examples will be discussed in section 6.

5.2 Implementation and comparison between PCT and ct-PCT

In the domain of food-borne diseases, the appropriate parsimony criterion is that of single-disorder explanations. This is an extreme form of the minimum cardinality criterion, which requires that the explanations should contain only one disorder. This criterion was implemented as a filter that runs after the programs have computed all irredundant explanations. Furthermore in order to collect information about the precision of ct-PCT in comparison with PCT, we registered also the number of irredundant solutions each program computed before the filter was applied.

Both the ct-BIPARTITE (for ct-PCT) and the BIPARTITE (for PCT) algorithms were implemented in Prolog, almost straightforwardly from their definition.

We tested the systems with seven artificial (non-clinical) cases created by the specialists. The set of observations and negative observations, together with their temporal information, were given to the program as a list, so the function

`NextObservation` in `ct-BIPARTITE` (and `BIPARTITE`) would just read in the next element on the list. The same list, stripped of temporal information and negative observations was the input for the `BIPARTITE` program.

On average, the number of single-disorder diagnosis computed by `ct-BIPARTITE` was less than half of the number of single-disorder diagnosis computed by `BIPARTITE`, and considering irredundant explanations, `ct-BIPARTITE` generated less than one third of the irredundant solutions generated by `BIPARTITE`. In terms of execution time, on average, `ct-BIPARTITE` computed all irredundant explanation in 70% of the time it took `BIPARTITE` to process the same example. In a particular example, `ct-BIPARTITE` computed only a single one-disease diagnosis, and 2 multiple diseases diagnosis against 6 one-disease diagnosis and 73 multiple-diseases diagnosis computed by `BIPARTITE`. For this example, `ct-BIPARTITE` completed its execution in 30% of the time it took to `BIPARTITE` to complete.

6 Conclusions

This work has presented two extensions of the original Parsimonious Covering Theory. The first extension `t-PCT` allows one to associate to each disorder an evolution of manifestations, and the second `ct-PCT` adds categorical information about the necessity or possibility of a disorder causing a manifestation to the temporal reasoning. These two extensions include the original `PCT`, in the sense that if the knowledge base contains no temporal or categorical information then `t-PCT`, `ct-PCT` and `PCT` will compute the same solution for all diagnostic problems.

A diagnostic systems for food-borne diseases was developed. The experimental results showed that in that domain, and we expect in all temporally rich domains, the temporal and categorical information allow for a faster and more precise diagnosis than the standard `PCT`.

The temporal/categorical extension to `PCT` has some limitations in representing both the knowledge base and case information in some diagnostic domains. Overcoming these limitations are possible areas of future research:

- As discussed in section 3.5, temporal graphs do not allow for cycles. In medical diagnosis few but important diseases have recurrent events. Malaria is one of them: one distinguishes different forms of malaria by the period between the re-occurrence of the fever episodes [1].
- Also discussed in section 3.5, `ct-PCT` does not allow to state that an observation has already happened and is no longer present without stating explicitly the time and duration for the observation.

- The semantics of arcs in temporal graphs, and the DIST function, refer to the beginning of the manifestations, that is, $\text{DIST}_I((m_x, m_y)) = I$ means that the elapse time between the *beginning* of manifestation m_x and m_y should fall within the interval I . This semantics of associating the beginning of the manifestations was appropriate for most of the examples, but not for all. In the example of intoxication by *A. phalloides*, *A. virosa*, and *A. verna* mushrooms (figure 2), the hepatorenal failure will occur 1 to 2 days after the *end* of the cramps and diarrhea symptoms. In that case, because the duration of the cramps and nausea symptoms were given, we could determine an interval from the beginning of the corresponded symptoms.
- In the food-borne diseases domain we encountered an example of categorical information that could not be represented in ct-PCT: representing that a disorder necessarily does not cause a manifestation. Uncomplicated developments of verotoxigenic stains of *E. coli* infections will cause bloody diarrhea but not fever.
- PCT's incapacity to represent the interference of disorders may be particularly severe in ct-PCT. In the presence of temporal information it is very unlikely that two disorders will not interfere with each other. As an example, let us suppose that d_i causes m with duration I and d_j causes m with duration J . Then certainly the presence of both disorders simultaneously will cause some change on the duration of m (the same can be true for the temporal relation of m with other manifestations in both d_i and d_j). In the area of medical diagnosis, this has been documented by [11]. If either I and J are mutually inconsistent ($I \cap J = \emptyset$), or either one is inconsistent with the duration of the observation m ($\text{DUR}^+(m) \cap I \cap J = \emptyset$) then the hypothesis that contains both d_i and d_j will be discarded as temporally inconsistent with the case.

An interesting extension of the ct-PCT theory, suggested to us by the specialists, is the incorporation of fuzziness into the idea of temporal inconsistency. The specialists were willing to accept a temporally inconsistent disease in a diagnosis provided its temporal information would not “disagree much” with the case information. This indicates that temporal consistency should not be modeled as a boolean attribute, but as a fuzzy one. If a measure falls within the interval it is fully consistent, and its degree of consistency would decrease the further away it falls from that interval. The idea of fuzziness must then be carried over to all other concepts in the theory.

6.1 Acknowledgments

The first author was partially supported by CNPq grant 521464/93-5. The second author was supported by a CNPq scholarship. We are very grateful to Renato Sabbatini (Núcleo de Informática Biomédica da UNICAMP),

Cézar Laerte Natal (Faculdade de Ciências Médicas da Universidade Federal de Uberlândia), and Fernando Gonçales (Instituto de Moléstias Infecciosas do Hospital das Clinicas da UNICAMP) for discussions on medical diagnosis in general and infections diseases in particular. We are in debt to the anonymous reviewers for suggestions that much improved this work.

References

- [1] R. Berkow, editor. *The Merk Manual*. Merk & Co, Inc., 16th edition, 1992.
- [2] L. Console and P. Torasso. On the co-operation between abductive and temporal reasoning in medical diagnosis. *Artificial Intelligence in Medicine*, 3:291–311, 1991.
- [3] L. Console and P. Torasso. Temporal constraint satisfaction on causal models. *Information Sciences*, 68:1–32, 1993.
- [4] C. Dousson, P. Gaborit, and M. Ghallab. Situation recognition: representation and algorithms. In *Proc. of the 13th IJCAI*, pages 166–172, 1993.
- [5] I. Hamlet and J. Hunter. A representation of time for medical expert systems. In J. Fox, M. Fieschi, and R. Engelbrecht, editors, *Lecture Notes in Med. Informatics*, volume 33, pages 112–119. Springer-Verlag, 1987.
- [6] K. Konolige. Abduction versus closure in causal theories. *Artificial Intelligence*, 53:255–272, 1992.
- [7] L. Console, D.T. Dupré, and P. Torasso. A theory of diagnosis for incomplete causal models. In *Proceedings of the 10th IJCAI*, pages 1311–1317, 1989.
- [8] W. Long. Reasoning about state from causation and time in a medical domain. In *Proc. of the AAAI 83*, 1983.
- [9] P. Lucas. Modelling interactions for diagnosis. In *Proceedings of CESA'96 IMACS Multiconference: Symposium on Modelling, Analysis and Simulation*, volume 1, pages 541–546, 1996.
- [10] G. L. Mandell, R. G. Douglas, and J. E. Bennett, editors. *Principles and practice of infectious diseases*. Churchill Livingstone, 4rd edition, 1995.
- [11] R. S. Patil. Causal representation of patient illness for electrolyte and acid-base diagnosis. Technical report, MIT Laboratory for Computer Science, 1981.
- [12] Y. Peng and J. A. Reggia. *Abductive Inference Models for Diagnostic Problem-Solving*. Springer-Verlag, 1990.
- [13] D. Poole. Representing diagnostic knowledge for probabilistic Horn abduction. In *Proc of the 12th IJCAI*, pages 1129–1135, 1991.

- [14] J. Reggia, D. Nau, P. Wang, and Y. Peng. A formal model of diagnostic inference. *Information Sciences*, 37:227–285, 1985.
- [15] E. Shortliffe. *Computer Based Medical Consultation: MYCIN*. Elsevier, 1976.
- [16] P. Torasso and L. Console. *Diagnotic Problem Solving: Combining Heuristic Approximate and Causal Reasoning*. Van Nostrand Reinhold, New York, 1989.
- [17] S. Tuhrim, J. Reggia, and S. Goodal. An experimental study of criteria for hypothesis plausibility. *J. of Experimental and Theoretical Artificial Intelligence*, 3:129–144, 1991.
- [18] S. Weiss and C. Kulikowski. Glaucoma consultation by computers. *Comp. Biol. Med.*, 8:25–40, 1978.