A Fuzzy Temporal/Categorical Extension to the Parsimonious Covering Theory

Jacques Wainer
Instituto de Computação
UNICAMP, Brazil
wainer@dcc.unicamp.br

Sandra Sandri
LAC
INPE, Brazil
sandri@lac.inpe.br

Abstract

This paper proposes a way of incorporating fuzzy temporal reasoning within diagnostic reasoning. Disorders are described as an evolving set of necessary and possible manifestations. The expected duration of each manifestation and the interval of time expected to occur between the beginning of any two manifestations are modeled by fuzzy sets. The patient information about the beginning and duration of the manifestations is also described using fuzzy sets and different measures on how well a disorder explains the patient's manifestations are defined.

1 Introduction

Diagnostic problem solving has been an area of intense interest in Artificial Intelligence, and has generated many methodologies, theories and applications over the last two decades. Diagnostic systems vary from rule-based systems [2], set-based theories [12, 11], logic based theories [13, 9, 7], and case-based reasoning [8]. Furthermore, temporal reasoning within diagnosis has long been considered an important part of diagnostic reasoning in some domains [10, 6, 3].

This paper can be classified as a contribution in the area of temporal reasoning for set-based approaches to model-based diagnostics. It extends Parsimonious Covering Theory (PCT), which is the theoretical foundation to our approach to diagnostics, so that it can deal with temporal measures modeled by fuzzy sets. This extension makes use of some of the mechanisms employed in systems dealing with fuzzy temporal information [5, 15, 16, 1].

The paper plan is as follows. Section 2 briefly discusses PCT's and Section 3 discusses a temporal categorical extension to PCT's that is the basis of this work. Section 4 defines the global consistency index of a disorder, which describes how well a disorder explains a set of manifestations. Section 5 discusses some further information besides the global consistency index that can be obtained from the model, in particular, information about which symptoms are most interesting to look for.

2 Basics of Parsimonious Covering Theory

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, as opposed to approaches such as the rule based one, that mix all these different forms of knowledge.

The basic version of PCT [12], deals mainly with general knowledge about diagnostics and the domain knowledge. The domain knowledge is represented by a set of diseases D, a set of manifestations M and a causal relation $C \subseteq D \times M$. A pair $\langle d_i, m_j \rangle \in C$ represents the fact that the disease d_i may cause the symptom m_j. This knowledge is usually easy to obtain.

The set of all symptoms that $d \in D$ may cause is denoted by $effects(d)$, and is defined as $\{m_i | \langle d, m_i \rangle \in C\}$. The set of all symptoms that the disorders in $D_L \subseteq D$ may cause is given by $effects(D_L) = \bigcup_{d \in D_L} effects(d)$. A $D_L \subseteq D$ is said to be a **cover** of a set of manifestations $M_l \subseteq M$ iff $effects(D_L) \supset M_l$. Similarly, we define all the possible causes of a manifestation $m \in M$ as $causes(m) = \{d_i | \langle d_i, m \rangle \in C\}$. This definition is also extended for sets of manifestations.

The information about a particular case is described by a set of manifestations M^+ that a patient presents during a given period of time.

A particular diagnostic problem is then defined by the knowledge base described by sets D, M and C, and by the factual information described in M^+. An **explanation** or a **diagnostic** for the case is a set $D_L \subseteq D$ such that $D_L \supset M^+$ and D_L satisfies some given parsimony criteria. The usual parsimony criteria are:

- 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.

Finally, the solution of a diagnostic problem is the set of all explanations for the case.

Although in many diagnostic application the appropriate parsimony criterium is of the minimum cover, [12] affirms that for computational reasons and for the sake of generality the irredundancy criterium is more interesting: from the set of all irredundant explanations one can algorithmically select all minimum explanations, and given the information in C one can also algorithmically generate all relevant explanation from the irredundant ones.

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. 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 to not interfere with each other. It is not possible to represent, in the original PCT, 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 [11].

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 [14] confirm. Domain specific heuristics are needed to filter the solutions obtained by the PCT so that only the most "appropriate," or "plausible," or "interesting" diagnostics are presented, or at least those diagnostics are presented first to the people using the system.

3 Overview of Temporal PCT and Categorical PCT

In [17] an extension to PCT was proposed so that categorical and temporal information could be added to the knowledge base. Temporal information allows one to represent the expected evolution of the manifestations caused by each disorder. Categorical information distinguishes between manifestations that are only possible from manifestations that are necessary in the course of a disease. Thus the knowledge about a disease may specify that it 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. Furthermore, it may specify that both m_1 and m_2 are necessary manifestations of a given disease, whereas m_3 is only a possible manifestation of that disease.

3.1 Temporal PCT

The representation of temporal information is accomplished by associating a graph $G_l =< V_l, A_l >$ to each disorder d_l, in which the nodes in V_l represent manifestations and the directed arcs in A_l represent temporal precedence. Quantitative information about the duration of a given manifestation is associated with its corresponding node, and quantitative information about the elapsed time between the beginnings of any two manifestations is associated with the corresponding arc. Quantitative information is modeled by intervals $I = [I^-, I^+]$ defined in a given time scale; this feature allows for the management of the imprecision inherently present in medical diagnosis, which wouldn't be accomplished if only single numbers where employed. Figure 1 illustrates the graph associated with disorder d_8.

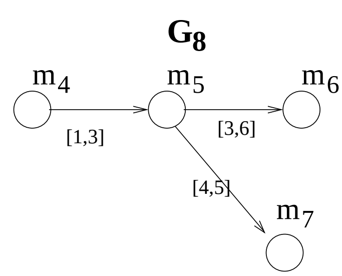

Figure 1: Example of a graph associated to a disorder

A similar representation is used for the case information. It is possible to state that, for a particular patient, manifestation m_4 started sometime between 2 and 3 weeks ago and lasted for 1 to 2 days, that manifestation m_5 started 5 days ago and is still present at the moment of consultation, and also that the patient presents manifestation m_6 but cannot state when it started.

Comparing the actual beginning of a manifestation in M^+ and its expected beginning in relation to a disorder d, as well as the expected and actual duration of those manifestations, one is capable of reducing the set of possible explanations to the case at hand. For instance, let us suppose that for a given disorder d_1, manifestation m_4 was supposed to start 1 week ago. Knowing that m_4 actually started between 2 and 3 weeks ago makes the hypothesis of the occurrence of d be inconsistent with the case data, and therefore any explanation containing d can be discarded for the case at hand. A similar reasoning is used in relation to the actual and expected duration of manifestations.

3.2 Categorical Temporal PCT

The representation of categorical information is accomplished by attaching either a P or a N label to each node in the graph, indicating whether that manifestation is necessarily or only possibly caused by the disease.

Information about a case in a categorical PCT is modeled by M^+, as described previously, and by M^-, which is the set of manifestations that are known to be absent in that case.

The categorical information allows one to make intelligent use of information about manifestations that are not present in the case at hand. Indeed, if a necessary manifestation is known to be absent, and the expected time of its beginning in relation to a given disease has already elapsed, then that disease cannot be part of any explanation. This feature thus allows one to further decrease the number of possible explanations to the case at hand.

4 Fuzzy Temporal/Categorical PCT

In this paper we will extend some of the ideas of the temporal PCT model in order to incorporate fuzzy concepts. However, we will make some simplification assumptions in relation to the original temporal model.

First of all, here we are only interested to verify how well a single disorder explains a set of manifestations, instead of constructing the whole diagnostic theory.

Moreover, we will assume that the temporal graph of disease d_l is a tree, and m_{0_l} is its root. That is not so strong an assumption as it may seem at first glance. In many diagnostic applications a disorder starts with an event, such as the breakdown of a component in mechanical or electrical systems, or infection of a patient by pathogens. It is this starting event that is modeled by the root "manifestation" of the temporal graph. Even though that root manifestation may not be observable, for modeling purposes one may assume that it is there.

Using this fuzzy extension, we will be capable of modeling more closely the kind of information furnished by both medical doctors and patients. Moreover, we will be capable of not only discarding diseases as before, but also ranking the more plausible ones. Finally, using these predictions we can furthermore rank the manifestations that should be priorily investigated, among those for which neither positive or negative actual information is available.

4.1 Basic Definitions

The knowledge base for a fuzzy temporal diagnostic problem is the information about a particular disorder and how it evolves. The knowledge base is given by $\langle D, M, G, \Theta, DIST, DUR, POSS \rangle$ where D is the set of disorders, M the set of manifestations, G is a function that associates to each disorder d_l a tree $G_l = \langle V_l, A_l \rangle$ that describes the development of the disorder d_l, where V_l are manifestations, and an arc from m_i to m_j states that the beginning of m_j occurs after the beginning of m_i. Θ is a time scale, and $DIST$ is a function that associates to each arc in each tree G_l a fuzzy temporal interval $DIST_l(m_i, m_j) = R$ [1] which 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 fuzzy temporal interval R. The duration function DUR associates with each node $m_i \in V_l$ in each graph G_l a fuzzy interval J, that specifies that the duration of m_i must be within interval J. Finally the function $POSS$ associates to each node $m_i \in V_l$ in each tree G_l either the label N or P, if the manifestation is necessary or possible in the development of the disorder d_l.

Information about a given case is modeled by a tuple $Ca = \langle M^+, M^-, \mathrm{BEG}^+, DUR^+, \theta_0 \rangle$. M^+ is the set of manifestations known to be or to have been present in the case. M^- is the set of manifestations known to be absent from the case. BEG^+ is a function that associates to each $m \in M^+$ a fuzzy temporal interval that represents the possible moments in which manifestation m started. DUR^+ is a function that associates to each $m \in M^+$ a fuzzy temporal interval which represents the possible duration of the manifestation. Finally θ_0 is the moment of the diagnosis.

Throughout this paper, we shall make use of 3 particular fuzzy intervals. Let θ_0 be the present moment of consultation. The fuzzy intervals describing the pos-

[1] A fuzzy interval R defined in Θ is characterized by a convex membership function $\mu_R : \Theta \to [0, 1]$, such that $\exists x \in \Theta, \mu_R(x) = 1$.

sibility of an event occurring *at any time*, *after the present moment*, and *before the present moment* are respectively defined as

- $I_{\text{anytime}} = A$, such that $\forall x \in \Theta, \mu_A(x) = 1$,
- $I_{\text{afternow}} = B$, such that $\forall x \in \Theta$, if $x \geq \theta_0, \mu_B(x) = 1$, and $\mu_B(x) = 0$, otherwise,
- $I_{\text{beforenow}} = C$, such that $\forall x \in \Theta$, if $x \leq \theta_0, \mu_C(x) = 1$, and $\mu_C(x) = 0$, otherwise.

We also make use of functions $\oplus$ and $\ominus$, which respectively yield the sum and subtraction of fuzzy sets. Let A and B be fuzzy sets, $A \oplus B$ and $A \ominus B$ are respectively characterized by membership functions [4]

- $\mu_{A \oplus B}(z) = sup_{\{(x,y)/z=x+y\}} min(\mu_A(x), \mu_B(y))$,
- $\mu_{A \ominus B}(z) = sup_{\{(x,y)/z=x-y\}} min(\mu_A(x), \mu_B(y))$.

Finally, we make use of function h, which for a given fuzzy set A in X, yields its height:

- $h(A) = \sup_{x \in X} \mu_A(x)$.

4.2 Covering

Our goal is to evaluate when and how much the disorder d_l explains the set of manifestations M^+. This means that the disorder must explain all the manifestations in M^+, and in the tradition of PCT must be a cover for M^+. Thus our first requirement is that

- $V_l \supseteq M^+$.

Clearly since, from the PCT perspective, an explanation contains only a single disorder, it satisfies all parsimony criteria.

4.3 Temporal/categorical consistency

The way the model of a disease predicts the evolution of the manifestations must be consistent with the way the symptoms are indeed evolving in a given case. If m_1 should occur before m_2 according to the disorder model and m_2 in fact occurred before m_1 in the case at hand, one can say that that disorder is not temporally consistent with the case. Similarly, if m_2 should occur 3 to 5 days after m_1 occurred, according to the model of a given disorder, and in the case at hand m_2 occurred 2 days after m_1, then again one can say that the disorder is temporally inconsistent with the case. However, in the later case one is less inclined on being too strict. The use of fuzzy concepts in evaluating the temporal consistency will allow us to discard the disorder in the first case, and still consider somewhat plausible the disorder in the later case.

In temporal consistency, we are not concerned if a manifestation is necessary or possible in relation to a given disorder. Categorical inconsistency, on the other hand, refers to the fact that a necessary manifestation of a disease must happen, provided there has been enough time for it to happen.

For instance, let us suppose that a patient makes a consultation on July, 10th, and tells the doctor that he started suffering from manifestation m_5 about 10 days ago, i.e. around July, 1st. Let us suppose the doctor, suspecting that the disorder affecting the patient could be d_8 (see Figure 1), asks the patient about the occurrence of manifestation m_6, to what the patient answers that he has not suffered from m_6 so far. Since the expected delay of time between m_5 and m_6 in d_8 is of about 3 to 6 days, then, supposing the patient was suffering from disorder d_8, m_6 should have already started between July 4th and July 7th. Therefore, even if the manifestation m_6 would start just after the consultation, the hypothesis of d_8 would be still inconsistent with the data. If m_6 is only a possible manifestation of d_8 then this inconsistency can be disregarded, since the patient does not necessarily have to suffer from m_6 given that he has d_8. However, if m_6 is a necessary manifestation in d_8, then this inconsistency allows the doctor to discard d_8 as an explanation for the patient's symptoms.

It should be noted that in temporal consistency, we only take the information of $M+$ into account, whereas in categorical inconsistency, we are also concerned with the information in M^-. Although we will incorporate temporal and categorical consistency in a single index, we treat these subjects separately below, for the sake of clarity.

4.3.1 Temporal consistency

Temporal consistency of a disorder d_l will be evaluated by figuring out, given the information of the starting moments for each manifestation present in M^+, when manifestation m_0 should have started. If there is not a possible moment for it to have started, then the case information is not consistent with the disorder model, and d_l can be discarded as a possible explanation.

Given M^+ and BEG^+, we will define the temporal consistency of a disorder d_l, denoted by $\alpha(d_l)$, as the height of the (backward) revised begin time of the root node. Namely,

- $\alpha(d_l) = h(\text{BEG}^+_{\leftarrow}(m_0))$, where $\mathrm{BEG}^+_{\leftarrow}(m_i)$ is defined below.

Before proceeding, let us establish the value $\mathrm{BEG}^+(m_i)$ for $m_i \notin M^+$:

- $\forall m_i \notin M^+, \mathrm{BEG}^+(m_i) = I_{\text{anytime}}$.

Therefore, for manifestations that have not occurred, or whose occurrence or not occurrence is not known, we suppose that their beginnings can happen (or could have happened) at any moment of the time scale.

Let L_l be the set of leaf nodes in G_l. The revised begin time of a manifestation m_i is defined as:

- $\forall m_i \in L_l, \mathrm{BEG}^+_{\leftarrow}(m_i) = \mathrm{BEG}^+(m_i)$.
- $\forall m_i \notin L_l$, $\mathrm{BEG}^+_{\leftarrow}(m_i) = \mathrm{BEG}^+(m_i) \cap \bigcap\{\mathrm{BEG}^+_{\leftarrow}(m_j) \ominus DIST(m_i, m_j)/(m_i, m_j) \in A_l\}$.

Intuitively, $\mathrm{BEG}^+_{\leftarrow}$ of a manifestation is the overlap between the period of time bounding its actual beginning ($\mathrm{BEG}^+(m_i)$), and the period of time bounding its expected beginning in d_l.

In order to determine the expected beginning of m_i in relation to d_l, we first determine when m_i should have begun, according to each of the manifestations m_j which are supposed to happen after m_i in d_l. For each of these m_j, this value is determined by taking into consideration the actual beginning of m_j and the elapsed time the model predicted between the occurrence of m_i and m_j ($\mathrm{BEG}^+_{\leftarrow}(m_j) \ominus DIST(m_i, m_j)$). Finally, we take the overlap of these values, since the restriction imposed by each m_j has to be satisfied ($\bigcap_{mj}(\mathrm{BEG}^+_{\leftarrow}(m_j) \ominus DIST(m_i, m_j))$.

4.3.2 Categorical consistency

Categorical inconsistency refers to the fact that a necessary manifestation of a disease must happen, provided there has been enough time for it to happen. Therefore, if at the moment of consultation there has been enough time of a necessary manifestation to happen and it has not yet occurred, then one can reject that disorder as the explanation for the manifestations.

One can say that a manifestation m_i had had enough time to occur in d_l if

- there exists a manifestation m_j, which was supposed to happen after m_i, i.e. $(m_i, m_j) \in A_l$, that has already occurred;
- or there exists a manifestation m_j, which was supposed to happen before m_i, i.e. $(m_j, m_i) \in A_l$, that did happen as expected, but the expected elapse time between the two manifestations has already expired.

The categorical consistency of a manifestation m in a disease d_l can be incorporated into the temporal consistency index. It is enough to initialize $\mathrm{BEG}^+(m)$ for $m \notin M^+$ as stated below and then calculate $\mathrm{BEG}^+_{\leftarrow}$ as defined previously. Let us define $N(d_l)$ as the set of necessary manifestations of d_l, $N(d_l) = \{m_i/POSS(m_i) = N\}$, then

- $\forall m_i \in M^- \cap N(d_l), \mathrm{BEG}^+(m) = I_{\text{afternow}}$,
- $\forall m_i \notin M^+ \cup (M^- \cap N(d_l)), \mathrm{BEG}^+(m_i) = I_{\text{anytime}}$.

Therefore, since $\mathrm{BEG}^+_{\leftarrow}(m_0)$ will now take into account also categorical information, index $\alpha(d_l)$ defined previously is said to be a temporal/categorical index.

4.4 Duration Consistency

The temporal information about the manifestations in a given disorder also includes information about the possible durations of each manifestation. Thus if the disorder model states that m_1 should last from 3 to 6 days and it did last around 7 days, one would like to say that the disorder is not completely consistent with the case regarding its duration specification.

The consistency of the duration of a manifestation m_i in relation to a disorder d_l is quantified as follows.

- For every node $m_i \in M^+$ such that $\mathrm{DUR}^+(m_i)$ exists, i.e. the manifestation has already ended, $\beta(m_i) = h(\mathrm{DUR}(m_i) \cap \mathrm{DUR}^+(m_i))$,
- For every node $m_i \in M^+$ such that $\mathrm{DUR}^+(m_i)$ does not yet exist, i.e. the manifestation has not ended yet, $\beta(m_i) = h(\mathrm{DUR}(m_i) \cap (I_{\text{afternow}} \ominus \mathrm{BEG}^+(m_i)))$.

A manifestation m_i that has already ended will be considered the more consistent with d_l in what regards duration, the more possible are the values in the overlap between its actual duration ($\mathrm{DUR}^+(m_i)$), and its expected duration in relation to d_l ($\mathrm{DUR}(m_i)$).

If a manifestation m has started at the (fuzzily) bound interval of time $\mathrm{BEG}^+(m_i)$ but has not yet finished at the moment of the consultation, we first construct the hypothetical duration of m_i in d_l ($I_{\text{afternow}} \ominus \mathrm{BEG}^+(m_i)$) and then take its intersection with the its expected duration in relation to d_l ($\mathrm{DUR}(m_i)$).

After obtaining the duration consistency of each manifestation m_i in d_l, we aggregate these values to obtain the duration consistency of disorder d_l itself. This global index is defined as

- $\beta(d_l) = \inf_{m_i \in V_l} \beta(m_i)$.

4.5 Intensity Consistency

In some diseases, it is important to quantify the intensity with which some of its manifestations occur. For instance, let us suppose a given disorder is characterized by strong fever at some time during its development; in this case, it is reasonable to suppose that that disorder will be the less plausible, the lower the temperature of the patient.

In order to provide information about the intensity of manifestations in relation to disorders, the knowledge base contains a function INT, which attributes to each node m of each temporal graph G_l a fuzzy set $\mathrm{INT}(m)$ describing the intensity with which that manifestation is expected to occur in d_l. Each fuzzy set $\mathrm{INT}(m)$ is defined on its particular domain $\Omega_{\mathrm{INT}}(m)$.

For manifestations m for which intensity is not a relevant matter in d_l, $\mathrm{INT}(m)$ is constructed as $\forall x \in \Omega_{\mathrm{INT}}(m)$, $\mu_{\mathrm{INT}(m)}(x) = 1$. When the intensity can be quantified by a precise constant x^* in X, then we construct $\mathrm{INT}(m)$ as $\mu_{\mathrm{INT}(m)}(x) = 1$, if $x = x^*$, $\mu_{\mathrm{INT}(m)}(x) = 0$, otherwise.

In the same way, in order to provide information about the intensity of manifestations presented by a given patient, the case information contains a function INT^+, which attributes to each node $m \in M^+$ a fuzzy set $\mathrm{INT}^+(m)$ describing the intensity with which that manifestation occurred. Each fuzzy set $\mathrm{INT}^+(m)$ is defined on domain $\Omega_{\mathrm{INT}}(m)$.

The consistency of the intensity of a manifestation m, in relation to a disorder d_l, is quantified as follows:

- $\gamma(m) = h(\mathrm{INT}(m) \cap \mathrm{INT}^+(m))$.

Finally, for a disorder d_l its intensity consistency is given by

- $\gamma(d_l) = \inf_{m \in V_l} \gamma(m)$.

4.6 Global Temporal Consistency

For a given disorder d_l, we have thus obtained a temporal/categorical consistency index $\alpha(d_l)$, which measures how consistent are the beginnings of the manifestations in M^+ in relation to the model of d_l and the information in M^-, a duration consistency index $\beta(d_l)$, which measures how consistent are the durations of the manifestations in M^+ in relation to the model of d_l, and an intensity consistency index $\gamma(d_l)$, which measures how consistent are the intensity of the manifestations in M^+ in relation to the model of d_l.

We summarize these indexes into a single global consistency index as

- $\delta(d_l) = min(\alpha(d_l), \beta(d_l), \gamma(d_l))$.

This global consistency index can be used as one of the aspects in ranking different explanations for a single case. Different disorders will have different global consistency indexes and in the lack of other factors, disorders with higher indexes are better explanations for the case.

5 Evaluation of the consistency of unknown manifestations

Besides the global consistency index that states how well a disorder explains the set of manifestations M^+, one can derive some further information about that diagnostic. For instance, it may be important in an interactive system to determine which further information one should get from the case, in particular which further manifestations must be investigated in the system being diagnosticated.

In the PCT tradition, usually one has many possible diagnostics, and thus it is more useful to prove that a particular disorder cannot explain a case than to accumulate evidence in favor of a particular diagnosis.

In order to evaluate when a manifestation should occur (or should have occurred) we propose that the expected beginning time $\mathrm{BEG}^+_{\leftarrow}$ be evaluated from leaf nodes towards the root, as it was done in section 4.3, but then also propagate that information from the root towards the leaves, generating the final expected beginning time $\mathrm{BEG}^+_{\rightarrow}$. This interval will contain the best information available on when a particular manifestation in $V_l - (M^+ \cup M^-)$ should occur or should have occurred. Finally, the intersection of that interval with $I_{\mathrm{beforenow}}$ defines how much of that interval lies before the moment of diagnosis. If there is no such intersection then the manifestation can only happen in the future. Thus trying to obtain information about this manifestation is interesting in order to eliminate the disorder as a possible explanation for the manifestation: if the manifestation did occur, it would be temporally inconsistent with the disorder.

In order to calculate $\mathrm{BEG}^+_{\rightarrow}$ we use $\mathrm{BEG}^+_{\leftarrow}$ as defined above and then that information is propagated forward:

- $\mathrm{BEG}^+_{\rightarrow}(m_0) = \mathrm{BEG}^+_{\leftarrow}(m_0)$,
- $\forall m_j \neq m_0, \mathrm{BEG}^+_{\rightarrow}(m_j) = \mathrm{BEG}^+_{\leftarrow}(m_j) \cap (\bigcap\{\mathrm{BEG}^+_{\rightarrow}(m_i) \oplus DIST(m_i, m_j) / (m_i, m_j) \in A_l\}$.

Comparing the actual time of consultation with the values $\mathrm{BEG}^+_{\rightarrow}(m)$ for a manifestation $m \in V - (M^+ \cup$

M^-), we can obtain a degree of relevance of trying to obtain information about that particular manifestation in relation to other ones in $V - (M^+ \cup M^-)$, in the context of disorder d_l:

- $\rho_l(m) = 1 - h(\text{BEG}^+_{\rightarrow}(m) \cap I_{\text{beforenow}})$.

Measure $\rho_l(m)$ increases as the updated expected beginning of m occurs after the moment of consultation. Therefore, if the doctor finds out that this manifestation has already occurred, the disorder at hand can be eliminated as a possible explanation of the manifestations. Indeed, if $\rho_l(m) = 1$ and m has already occurred, then running the system again with $m \in M^+$ will yield a temporal inconsistency.

Given a set of disorders D, we can obtain the global relevance of questioning about a doubtful manifestation m by making

- $\rho(m) = \sup_{\{d_l \in D\}} min(\rho_l(m), \gamma(d_l))$

Therefore, the questions with higher priority will be the ones that score a high priority in relation to disorders which themselves have a high possibility of being the good explanation to the patient's symptoms. This feature allows us not only ranking the most important questions to ask, but also to eliminate completely irrelevant questions, i.e. about manifestations that are only meaningful in disorders that cannot explain the case data anyways.

6 Conclusions and future work

This paper has presented some first steps towards incorporating fuzzy temporal reasoning into PCT's. It can be seen as an extension to the model presented in [17], by modeling temporal data as fuzzy sets, and by addressing the intensity with which manifestations can appear in a given disorder (also using fuzzy sets). As a consequence of the use of fuzzy sets, it is possible to order the the disorders most likely to explain a set the manifestations. The approach presented here, however, is restricted to disorder models that can be structured in a tree, instead of any direct acyclic graph as in [17].

As in the approaches found in [15, 16, 1], here the temporal pieces of information are modeled by fuzzy sets and the usual fuzzy operators are used to manipulate them. This work is not concerned with the internal consistency of the data, e.g. the whole corpus of information furnished by a patient in a medical case, as found in [15], or with the language used to obtain them, as found in [1]. It has been conceived to be specifically used in diagnosis and, for this reason, focuses more on the disorder models than on the case data as in those approaches. In this sense, it is more related with [16] but, unlike that work, does not use a logical approach to the problem.

The approach presented here also distinguishes between incompatibility between expected and factual durations of manifestations and incompatibility between expected and factual beginning of manifestations, whereas the other approaches merge together these two aspects of temporal information. The approach presented here yields only possibilistic compatibility degrees, but could be easily modified to obtain also entailment degrees, as in those approaches. Entailment degrees such as necessity measures, although considered to be too restrictive by the authors, may be useful to distinguish between two diagnostics yielding the same possibilistic degree of compatibility in a given case. Finally, this work incorporates other important aspects of diagnosis as assessing the compatibility of the intensity of manifestations and ordering manifestations that are more likely to happen in a given context.

We are currently working in expanding the results obtained here in some different directions:

- allowing the graph associated with each disorder to be a multi-source, directed, acyclic graph, instead of a tree. This has consequences in calculating both the backward and forward revised beginning time $\text{BEG}^+_{\leftarrow}$ and $\text{BEG}^+_{\rightarrow}$, which become dependent of the path.
- investigating a different operation for the calculation of $\text{BEG}^+_{\leftarrow}$. In the expression $\text{BEG}^+_{\leftarrow}(m_j) \ominus DIST(m_i, m_j)$ used in the calculation, what is really needed is the inverse of the $\oplus$ operation, when it is defined [4]. The consequences of using the "inverse of $\oplus$" are narrower fuzzy intervals and thus more precision. But it is important to notice that by using $\ominus$ one is being conservative: a disorder will never be considered temporally inconsistent when it really is not inconsistent, although the reverse could be true.
- in the item "further information about the disorder," we are investigating how to make predictions about future manifestations. In medical domains it is usual that $DIST(m_i, m_j)$ will be a wide interval, which means that even if one has precise information about $\text{BEG}^+(m_i)$ there will be less precise information about when m_j should begin. But given how the case has been developing one can make some evaluation on whether the disorder is progressing faster or slower than the "average" and take that into consideration in forecasting future manifestations.

- allowing for "fuzzy" categorical information, making it possible to model pieces of information such as "in disorder d_l, manifestation m_i is more likely to happen than not to happen".

Finally, it is important to reaffirm that the theory developed herein deals only with a single disorder explaining the set of manifestations in the case. There are complications to model explanations that contain more than one disorder and in which the disorders in the explanation have manifestations in common. In the case that all disorders in an explanation do not have common manifestations it seems that the theory above could be generalized by calculating the δ index for each disorder and attributing the consistency index to the explanation the minimum of the disorder's indexes. We are currently investigating such approach.

In the case that disorders have common manifestations, the problem is that the indexes above depend on which manifestations one is attributing to a particular disorder. For example, a manifestation m_4 may be causing one of the consistency indexes of disease d_1 (in which m_4 is a possible manifestation) to be very low, but if one could state that the m_4 in the case "is being caused" by a different disorder, say d_2, then the indexes for d_1 would be better, as would the indexes for explanation $\{d_1, d_2\}$. This means that one has to include in the basic PCT framework a concept of which disorder is causing which manifestation. Furthermore, one would have to deal with the combinatorial explosion problem of always adding new disorders to improve explanation indexes.

Acknowledgements

We would like to thank the comments of an anonymous referee which contributed to the improvement of this work. The first author was partially supported by the CNPq/PROTEM project LOGIA. The second author was partially supported by the CNPq grants 350286/93-0 and 301272/97-2.

References

[1] S. Barro, R. Marin, J. Mira, and A.R. Paton. A model and a language for the fuzzy representation and handling of time. *FSS*, 61:153–175, 1994.

[2] B.G. Buchanan and E.H. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* Addison-Wesley Publishing Company, 1984.

[3] L. Console and P. Torasso. Temporal constraint satisfaction on causal models. *Information Sciences*, 68:1–32, 1993.

[4] D. Dubois and H. Prade. *Possibility Theory: an approach to computerized processing of uncertainty.* Plenum Press, 1988.

[5] D. Dubois and H. Prade. Processing fuzzy temporal knowledge. *IEEE Trans. on S.M.C.*, 19(4), 1989.

[6] 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.

[7] K. Konolige. Adbuction versus clusure in causal theories. *Artificial Intelligence*, 53:255–272, 1992.

[8] P. A. Koton. *Using Experience in Learning and Problem Solving.* PhD thesis, MIT Laboratory of Computer Science, MIT/LCS/TR441, 1988.

[9] 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.

[10] W. Long. Reasoning about state from causation and time in a medical domain. In *Proc. of the AAAI 83*, 1983.

[11] 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.

[12] Y. Peng and J. A. Reggia. *Abductive Inference Models for Diagnostic Problem-Solving.* Springer-Verlag, 1990.

[13] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, April 1987.

[14] 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.

[15] L. Vila and L. Godo. On fuzzy temporal constraint networks. *Mathware and Soft computing*, 3:315–334, 1994.

[16] L. Vila and L. Godo. Possibilistic temporal reasoning on fuzzy temporal constraints. In *Proc. IJCAI'95*, 1995.

[17] J. Wainer and A. Rezende. A temporal extension to the parsimonious covering theory. *Artificial Intelligence in Medicine*, 10:235–255, 1997.