Automated Trend Detection with Alternate Temporal Hypotheses

Ira J. Haimowitz
MIT Laboratory for Computer Science
545 Technology Square, Room 414
Cambridge, MA 02139
ira@medg.lcs.mit.edu

Isaac S. Kohane
Children's Hospital, Harvard Medical School
3(X) Longwood Avenue
Boston, MA 02115
gasp@medg.lcs.mit.edu

Abstract

We have written a prototype computer program called $TrenD_x$ for automated trend detection during process monitoring. The program uses a representation called *trend templates* that define disorders as typical patterns of relevant variables. These patterns consist of a partially ordered set of temporal intervals with uncertain endpoints. Bound to each temporal interval arc value constraints on real-valued functions of measurable parameters.

$TrenD_x$ has been used to diagnose trends in growth patterns from examining heights, weights and other parameters of pediatric patients. As $TrenD_x$ analyzes successive data points, the program updates its hypotheses about which stage of the growth process each data point belongs to. We present an example of $TrenD_x$ reaching temporally plausible diagnoses for an actual patient with delayed growth currently being seen at Boston Children's Hospital.[1]

1 Introduction

This work is part of the growing body of artificial intelligence (AI) research on diagnostic process monitoring. We wish to automatically detect trends, defined thusly:

- A *trend* is a clinically significant pattern in a sequence of time-ordered data.

These trends may be multivariate, and may consist of several distinct phases. Our trend detection program, called $TrenD_x$, can identify a trend and give a chronology of which data were in each phase.

We are particularly motivated by application domains where there is no reliable structure-function model because the underlying mechanism is poorly understood. Our application domain, pediatric growth, is such an area. Physicians recognize many hormonal, nutritional and genetic factors for growth, but cannot predict quantitative effects on height or weight from changes to any of these factors. Therefore in this domain the recent AI work in monitoring using semi-qualitative simulation [Dvorak and Kuipers 1989) or Bayesian networks [Berzuini, Bellazzi et al. 1992) may be insufficient.

Time-series analysis techniques [Avent and Charlton 1990], when used with a curve-fitting model of process disorders, are potentially useful for trend detection as we have described it. However, most curve-fitting models of pediatric growth [e.g. Thissen and Bock 1990] do not accurately match data of most new individuals and also contain many parameters that cannot be interpreted biologically. Furthermore, statistical models usually describe patterns after a fixed time point, while our goal includes detecting trends that may occur at any point in time.

Even in applications characterized by incomplete understanding of mechanism, domain experts can accurately detect trends. Furthermore they can verbally describe prototypical trends consisting of constraints that restrict certain values of variables over time.

We have written a prototype program called $TrenD_x$ for automated trend detection. The program uses a representation called *trend templates* that define disorders as typical patterns of relevant variables. These patterns consist of a partially ordered set of temporal intervals with uncertain endpoints. Bound to each temporal interval are value constraints on real-valued functions of measurable parameters. As $TrenD_x$ interprets data points of a process, the flexible temporal constraints allow alternate hypotheses of how that process has varied over time.

2 Pediatric Growth Monitoring

Our initial application domain for automated trend detection is pediatric growth monitoring. The principal tool pediatricians use to monitor the growth of their patients is the growth chart. Figure 1 shows the height of a child with age. A set of curves representing standard deviations (-2, -1, +1, +2) and the mean for heights of male children studied by the National Center for Health Statistics (NCHS) [Hamil, Drizd et al. 1979] are pre-plotted on the chart. Each standard deviation (SD) curve describes the proportion of the male children in the U.S.A. of the same age who arc taller or shorter than children whose height falls on that curve. For

1. This work has been supported (in part) by NIII grant R01 LM 04493, NICHHD 5T32 HD07277-9, and by a U. S. Office of Naval Research Graduate Fellowship.

instance, as calculated from the normal distribution, all children whose height is on the - 2 SD curve arc taller than 2.28% of male children in the U.S.A. of the same age. Pediatricians often use centiles rather than their corresponding standard deviations.

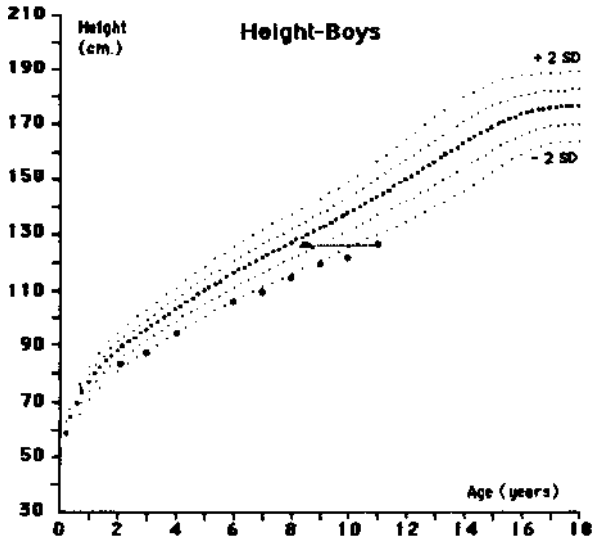

Figure 1 Growth chart height points of patient described throughout this article. Triangle shows hone age of X at chronological age of 11.

As a process monitoring application, pediatric growth is distinct in at least two ways. First, growth data is relatively clean; carefully measured heights are usually accurate within three millimeters. Second, although a child's height and weight are measured as infrequently as once per year, this sampling rate is frequent enough for expert pediatric endocrinologists to identify potential growth disorders. As we shall see, our program $TrenD_x$ does not require the monitored data to be sampled regularly or frequently.

2.1 An Example Case

The height points illustrated in Figure 1 are from the growth chart of a patient at the endocrinology clinic at Boston Children's Hospital. This patient, whom we shall refer to as Patient 002, was brought into the endocrine clinic at age 11 for consultation about possible growth disorders. The patient's general pediatrician had noticed a particularly sharp drop against the NCHS standards at age 10. This was of concern as children with a normal tempo of growth tend to grow on or parallel to the standard curves. The endocrinologists at Children's Hospital agreed that Patient 002 did not appear to exhibit an average growth pattern, and that one of the likely hypotheses was constitutional delay of growth, a nonnal variant of growth marked by delayed puberty and bone maturity well behind the patient's chronological age.

2.2 Clinically Significant Patterns:

From consulting clinical literature and experts in pediatric endocrinology, we are acquiring text descriptions of patterns followed by children with normal and abnormal growth. Two of these patterns, related to Patient 002, are described here. The descriptions use centiles rather than standard deviations.

The text description of average nonnal growth in boys, before puberty, consists of four general constraints:

- From birth until age 2 - 3 years, the patient establishes his or her centiles for height and weight. During this time height and weight centiles should vary in the same way from their original centiles.
- From then until the onset of puberty, the patient stays close to the same centiles in height and weight, with respect to the population of children having puberty at the average age.
- From birth until puberty, bone age is approximately equal to chronological age
- Puberty begins between age 10 and age 15, and is measured by when the qualitative testicular stage of the boy changes from stage 1 to stage 2.

The text description of constitutional delay of growth in boys differs only in that puberty begins between age 12 and 16, that the bone age is delayed 1 to 4 years behind chronological age, and that the reference population for height centiles is children having delayed puberty rather than puberty at the average age. These alternate reference populations were presented in [Tanner and Davies 1985].

Notice that the text descriptions above are characterized by uncertainty both in the *value* of relevant variables, and in the *times* over which these values hold. For example, there is uncertainty in the time of the onset of puberty, and uncertainty of the age at which the first constraint ends and the second begins. There is uncertainty in a child staying "close to the same centiles" and in the bone age being "approximately equal to" chronological age. This uncertainty is due to experts characterizing a common pattern for a large number of patients with the same growth trend.

3 TVend Templates

We wish to represent trends that include the uncertainty mentioned above. Additionally, our representation must provide enough constraints to distinguish between competing trends. Our attempt to reach these goals is the *trend template*.

A *trend template* is a prototypical pattern of data for a process exhibiting a certain behavior. In Figure 2 is the Wend template for male average growth before puberty adapted from the text constraints in section 2.2. Time constraints are drawn horizontally and value constraints are drawn vertically, The next two sections define trend templates and illustrate with this example.

3.1 Temporal Constraints

The temporal component of a trend template includes *landmark points* and *intervals*. The landmark points represent significant events in the lifetime of the monitored process. Landmark points may linked with time ranges *(min max)* expressing the minimal and maximal times between them. The trend template for average nonnal growth includes three landmark points: *birth, puberty onset,*

and *growth stops. Puberty onset* occurs 10 to 15 years alter *birth,* and *growth stops* occurs 17 to 19 years after *birth.* The links between landmark points establish a conceptual lime scale of the part of the process life being modeled. In the diagram, *birth* is illustrated as at time 0 and a time scale is labeled in "years." These arbitrary labels are for illustration; a trend template requires no "zero point" and allows mixing of time units (hours, years, etc.) in temporal distances.

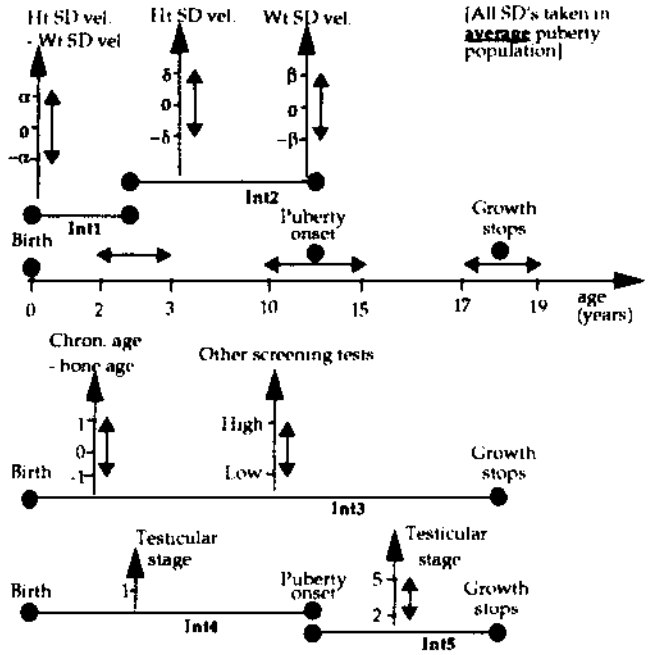

Figure 2 Trend template for average normal growth.

Intervals represent periods of the process that are significant for diagnosis or therapy. The above trend template consists of five intervals: *Intl* corresponds to infancy, *Intl* to mid-childhood, *Int3* to all of childhood, *Inl4* to pre-pubertal childhood, and *Int5* to post-pubertal childhood. Intervals consist of begin and end points whose times arc declared either as:

- offsets of the form *(min max)* from a landmark point, or
- offsets of the form *(min max)* from another interval begin or end point.

We represent time using the Temporal Utility Package (TUP) of [Kohane 1987J. TUP is a temporal reasoning program with both time points and time intervals; intervals include a begin point and an end point. Time is considered discrete, and for the growth domain $TrenD_x$ uses a minimal time distance of one second. With TUP, one asserts a variable temporal distance between two points in the form of a *range relation.* A TUP range relation has the general form:

```
(range p1 p2 :lb n1 :ub n2)
```

where p_1 and p_2 *are* points, and the n_1 and n_2 are integers denoting the lower and upper bounds on the numbers of seconds between p_1 and P_2. All temporal distances in a trend template are created with statements like these.

3.2 Value Coastraints

The value component of a trend template interval is a set of *value constraints,* each of which states that some function of a set of measurable parameters must fall within a certain range. Thus each value constraint is an expression of the form

$$m \leq f(D) \leq M \qquad \text{(EQI)}$$

where f is some real valued function defined on patient data, m is real-valued or $-\infty$, and M is a real-valued or $+\infty$. Each temporal interval in a trend template is associated with a set of value constraints. In the diagnostic program $TrenD_x$, the function f is evaluated on the set of data D currently assigned to that interval and the result is compared to the bounds m and M.

In the average normal growth template of Figure 2, interval *Intl* represents infancy, when height and weight centiles are established. We encode that height and weight centiles vary in the same way by constraining the difference between the average velocity of height SDs and the average velocity of weight SDs to be within a small number a of zero. Interval *Int2* represents the period of the boy staying in his centile channels, *Intl* begins at the endpoint of *Int I* and *Int2* ends at *puberty onset.* There are two value constraints: both the average velocities of height SDs and that of weight SDs are close to zero. Presently the small value constraint bounds in this trend template are $\alpha = \delta = \beta = 0.4$. These values may change as we refine the trend template.

Intervals *Int3* and *Int4* constrain other patient parameters. *Int3* begins at *birth,* ends at *growth stops* and describes the normal parameter values. Screening tests for growth disorders are within normal published values. Also, the difference between chronological age and bone age must be within one year of zero. *Int4* represents the pre-pubertal genital development: from *birth* until *puberty onset* by restricting testicular stage to 1. *Int5* distinguishes that puberty has occurred by constraining testicular stage to be between 2 and 5. Together these intervals encode the temporal and value constraints of our earlier text descriptions of male prepubertal average normal growth.

The trend template for male pre-pubertal constitutional delay of growth is similar to that of average normal growth. There arc a few differences. The time bounds on the landmark points *puberty onset* are 12 to 16 years, and the bounds on *growth stops* are 18 to 22 years. The value constraint in Int3 states that bone age is one to four years behind chronological age. Additionally, height SDs arc measured with respect to the population *of delayed puberty* patients, as opposed to those of average puberty in the average growth trend template.

4 Reasoning by $TrenD_x$

$TrenD_x$ diagnoses trends by matching process data to the constraints of trcnd templates. Those templates matching the data are retained as hypotheses; those not matching are discarded with possible triggering of other templates. For each disorder template, $TrenD_x$ may maintain multiple hypotheses with different chronologies of the data fitting the

intervals of the template. Here we describe the diagnostic algorithms and illustrate with the data of Patient 002.

Each $TrenD_x$ hypothesis hyp for a patient pat includes a trend template TT(hyp) and an assignment ASSIGNMENT(hyp) of patient data to the intervals of the template. ASSIGNMENT(hyp) is a relation [(int d)}, where int is an interval in TT(hyp) and d is an interpreted datum of pat. Each hypothesis also includes a set CONTEXT(hyp) (see next section) of temporal information in that assignment.

4.1 Initializing a Hypothesis for a Patient

In $TrenD_x$ every patient is initially assigned a hypothesis of average normal growth. The program does this using the context mechanism of TUP. A context is a collection of consistent temporal assertions that offers a single interpretation of events in the world - a possible chronology of events in the $TrenD_x$ hypothesis. Contexts are arranged hierarchically in a tree, with temporal assertions from a parent context also holding in child contexts. Every context has at most one parent.

The general architecture for contexts in $TrenD_x$ appears below:

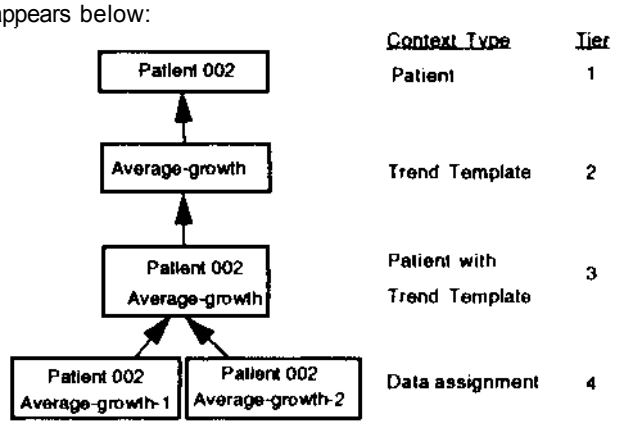

Figure 3 The $TrenD_x$ context hierarchy

The root context for each patient contains temporal information about that patient reported by the physician or hospital database, including the times of laboratory data and of significant life events. The birth date is asserted here. The times of events may be stated absolutely, by the Gregorian calendar, or relatively, within some bounds of another event in the patient context.

The average growth trend template presented earlier has its own context, which contains all of the points and range relations of that template. $TrcnD_x$ assigns a hypothesis to a patient by placing the context for the trend template of that hypothesis as a child of the patient context. We then place under the trend template context a third-tier context for temporal assertions about the patient's data matching the trend template. In this third-tier context we create a range relation that equates in time the birth date of the patient and the birth point of the trend template. With this assertion, TUP can calculate the temporal distance between any patient data point and any interval in the trend template. The hypothesis of average growth for Patient 002 initially only has the one chronology of the third tier context. When patient data are interpreted and assigned to intervals of the trend template, there may be alternate assignments and therefore alternate chronologies. In that case $TrenD_x$ branches to multiple hypotheses and associated (fourth tier) contexts.

4.2 Matching Algorithms

Let d be a patient datum, let hyp be a hypothesis of that patient, and let int be an interval in TT(hyp). We say that VALUE-SATISFIESW, int) iff all value constraints bound to int meet one of the following conditions:

1. The parameter of d is not constrained by that value constraint.
2. Inhere are insufficient data assigned to int to evaluate the function on the value constraint.
3. The function on the value constraint evaluated on d and the data assigned to int, is within the constraints range.

Ihe reasoning in $TrenD_x$ tits a data-driven process monitoring cycle. For each datum d in the input data stream, $TrenD_x$ executes the algorithm PROCESS-DATUMW).

Algorithm: PROCESS-DATUM [*d* : datum]
Let *pat* = the patient of *d*.
Add *d* to the list of data for *pat*.
For all hypotheses *hyp* constraining the parameter of *d*:
Let *matches* = MATCHING-INTERVALS(*d*, *hyp*)
(see algorithm below).
If *matches* = ∅
Remove *hyp* as a hypothesis for *pat*.
Else If length of *matches* = 1
Let *int* = the single interval in *matches*.
Add (*int d*) to ASSIGNMENT(*hyp*).
Else Let k = the number of temporally valid assignments $\{(int_1\ d)\ \ldots\ (int_j\ d)\}$ of *d* to subsets of intervals in *matches*.
Create k new hypotheses $hyp_1, \ldots, hyp_k$.
Create k new contexts $C_1, \ldots, C_k$ as children of CONTEXT(*hyp*).
For i = 1 to k
Let A_i = assignment i of the k assignments above.
Set CONTEXT(hyp_i) to C_i and set TT(hyp_i) to TT(*hyp*).
Add A_i to ASSIGNMENT(hyp_i).
Add hyp_i as a hypothesis of *pat*.
End For
End If
If no hypotheses remain for *pat*
Suggest a new hypothesis with a different trend template than TT(*hyp*).
End If
End For
END PROCESS-DATUM.

Algorithm MATCHING-INTERVALS takes a datum *d* and a hypothesis *hyp* and checks if *d* matches all necessary constraints of TT(*hyp*). If so, the algorithm returns the intervals in TT(*hyp*) that *d* matches. If not, ∅ is returned.

Algorithm: **MATCHING-INTERVALS**

[*d* : datum, *hyp* : hypothesis]

Let *relevant-ints* = the intervals in TT(*hyp*) that constrain the parameter of *d*.

[Note: We know that *relevant-intervals* is not empty from PROCESS-DATUM.]

Let *time-required-ints* =

{*int* ε *relevant-ints* | *int* must temporally include *d*}.

Let *time-possible-ints* =

{*int* ε *relevant-ints* | *int* may temporally include *d*, but not necessarily}.

[Note: *time-required-ints* ∩ *time-possible-ints* = ∅.]

Let *matched-time-possible-ints* =

{*int* ε *time-possible-ints* | VALUE-SATISFIES(*d*, *int*)}.

If ∀ *int* ε *relevant-ints*:

[*int* ε *time-required-ints*

⇒ VALUE-SATISFIES(*d*, *int*)]

Return *time-required-ints* ∪ *matched-time-possible-ints*

Else Return ∅.

End If

End MATCHING-INTERVALS.

The worst case complexity of PROCESS-DATUM is O(*k*), where *k* is the number of temporally valid assignments of *d* to the matching intervals in TT(*hyp*). In the worst case *d* can be temporally in any of five places (before, begin, between, end, after) compared to each trend template interval. Thus $k = O(5^I)$, where *I* is the number of intervals. However, in practice k is much smaller because the intervals represent process phases and tend to meet or be disjoint. Note that in the example *below k* never exceeds 3. This analysis assumes a unit time for all temporal queries in TUP. For the complexity of these queries consult [Kohane 1987].

4.3 TrenD$_x$ on the Example Case

Here we show how TrenD$_x$ applies these algorithms to the data of Patient 002. At first Patient 002 has one hypothesis; call it *hyp*. TT(*hyp*) is the male average growth trend template. ASSIGNMENT(*hyp*) = ∅.

The first datum hd_1 to be processed is a height at age 2.1. MATCHING-INTERVALS(hd_1, *hyp*) = {*Int1*, *Int2*}, two of the intervals in Figure 2. Recall the temporal relationship between these two intervals:

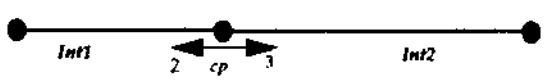

The end point of *Int1* and the begin point of *Int2* are equal in time. Thus we depict them as meeting at a common point *cp*, occurring between 2 and 3 years. Because hd_1 falls between 2 and 3 years, TrenD$_x$ concludes that there are three interval assignments, illustrated in Figure 4. The assignments are bound to the new hypotheses hyp_1, hyp_2 and hyp_3.

ASSIGNMENT(hyp_1) = {(*Int1*, hd_1)},

ASSIGNMENT(hyp_2) = {(*Int1*, hd_1), (*Int2*, hd_1)},

ASSIGNMENT(hyp_3) = {(*Int2*, hd_1)}.

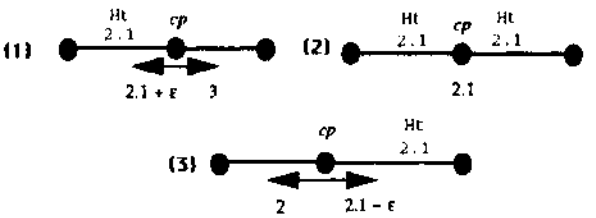

Figure 4 Alternate data assginments for the height at age 2.1 to intervals of the male average normal growth trend template. *Cp* is the common point shared by *Int1* and *Int2*.

TrenD$_x$ removes *hyp* from the hypotheses of Patient 002 and adds hyp_1, hyp_2, and hyp_3. TrenD$_x$ also adjusts the timing of *cp* in the context of each hypothesis, as shown above. Next TrenD$_x$ processes a weight datum wd_1, also at age 2.1, and assigns it to the same intervals as hd_1 in each hypothesis.

4.3.1 Pruning Hypotheses

Next TrenD$_x$ processes the second height datum hd_2 at age 3.0. The height SD on average pubertal standards for hd_1 is -1.25 and that of hd_2 (0.9 years later) is -2.52. The average velocity of -1.41 SDs per year is well below the value constraint on height of *Int2*. Therefore for hyp_2 and hyp_3 VALUE-SATISFIES(hd_2, *Int2*) is false. As shown in Figure 5, TrenD$_x$ prunes hyp_2 and hyp_3 from the hypotheses of Patient 002.

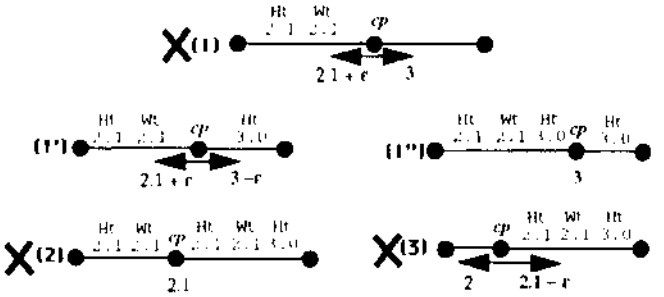

Figure 5 Pruning and branching of temporal hypotheses when adding the height at age 3.0 to the trend template for average growth.

TrenD$_x$ also branches to two alternate chronologies for hyp_1. Because the time constraint on *cp* in CONTEXT(hyp_1) is between 2.1+ ε and 3, hd_2 can be assigned to only *Int2* or to the common point of *Int1* and *Int2*. TrenDx creates two new hypotheses hyp_1' and hyp_1'', with associated contexts and assignments shown in Figure 5. Hyp_1 is removed from the hypotheses of Patient 002, while hyp_1' and hyp_1'' are added.

TrenD$_x$ continues to compare data of Patient 002 to hyp_1' and hyp_1''. The weight at age 3.0 results in pruning of hyp_1'' due to a failed value constraint in *Int1*. The only remaining hypothesis is hyp_1'. CONTEXT(hyp_1') contains the chronology that *cp* is somewhere between 2.1 and 3 years, noninclusive. Hyp_1' persists through several years of data as

TrenD$_x$ assigns more heights and weights to *Intl.* Later, Patient 002's height SD for the average pubertal population drops from -2.08 (age 6) to -2.61 (age 7). Call the height at age 7 hd_7. VALUE-SATISFIES(hd_7, *Int2)* is false and the final average growth hypothesis hyp_1 is pruned.

4.3.2 Triggering a New Trend Template

Once no average growth hypotheses remain for Patient 002, TrenD$_x$ considers the alternate hypothesis of delayed growth. Details of how TrcnD$_x$ triggers alternate disorders are presented in [Haimowitz and Kohane 1993J. The program monitors all of Patient 002*s growth data in order. TrenDx branches similarly with the delayed growth hypothesis as it had with average growth (see Figure 6). However, VALUE-SATISFIES(hd_7, *Intl)* is true for the delayed growth template because that template refers to the population of delayed pubertal boys. On this population standard the height SDs stay close to equal. *Hypj* • lasts three years longer, until age 10. It was at this age that the patient's pediatrician became concerned about a possible pathological growth disorder.

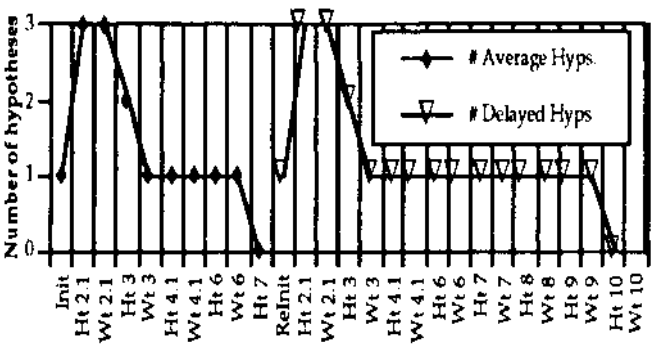

Figure 6 Number of hypotheses for growth patterns as TrenD$_x$ processes data of patient 002.

5 Related Literature

Different approaches to automated monitoring and trend detection were mentioned in section 1. Other researchers [e.g. Allen and Koomen 1983J have encoded temporal predicates in associative rules to test conditions in a process for a diagnostically significant pattern. TrenD$_x$ extends this research by representing the entire process as phases.

The different temporal interpretations of the same data in TrenD$_x$ is related to the Time Map Manager [Dean and McDermott 1987], that maintains different accounts of which logic propositions are true over which temporal intervals. TrenD$_x$ differs by interpreting primary numerical data.

6 Conclusions and Future Work

Trend templates represent multi-variatc trends as constraints on parameters over intervals that correspond to phases of a process. This representation is based on how expert diagnosticians verbally report their knowledge of trends. For this reason trend templates may be useful for knowledge acquisition and explanation of trends.

TrenD$_x$ monitors process data and matches them to hypotheses which include a trend template and a chronology of how the data fall into different stages of the trend. Our prototype application to growth chart monitoring produces plausible hypotheses on a real patient.

We will extend our trend template representation to detect trends that may occur at any time (initial work appears in [Haimowitz and Kohane 1993]). We are considering probabilistic bounds on value constraints, which may be used for assigning numerical scores to the match of data to template.

TrenD$_x$ could detect trends more flexibly by representing data measurement error, or by ignoring markedly aberrant data. TrenD$_x$ may be more useful by presenting a ranked differential diagnosis of several likely trends.

7 References

Allen, J. F and J. A. Koomen (1983). "Planning Using a Temporal World Model." *International Joint Conference on Artificial Intelligence (IJCAI-83),* pages *741-747.*

Avent, R. K. and J. D. Charlton (1990). "A Critical Review of Trend-Detection Methodologies for Biomedical Monitoring Systems." *Critical Reviews in Biomedical Engineering, 17(6): 621-659.*

Berzuini, C, R. Bellazzi, S. Quaglini and D. J. Spiegelhalter (1992). "Bayesian Networks for Patient Monitoring." *Artificial Intelligence in Medicine,* 4: 243-260.

Dean, T. L. and D. V. McDermott (1987). "Temporal Data Base Management." *Artificial Intelligence,* 32:1-55.

Dvorak, D. and B. Kuipers (1989). "Model-Based Monitoring of Dynamic Systems." *International Joint Conference on Artificial Intelligence (IJCAI-89),* pages 1238-1243.

Haimowitz, I. J. and I. S. Kohane (1993). "An Epistemology for Clinically Significant Trends." *National Conference on Artificial Intelligence (AAAI-93),* lo appear.

Hamil, P. V., T. A. Drizd, C. L. Johnson, R. B. Reed, A. F. Roche and W. M. Moore (1979). "Physical Growth: National Center for Health Statistics Percentiles." *The American Journal of Clinical Nutrition,* 32: 607-629.

Kohane, I. S. (1987). *Temporal Reasoning in Medical Expert Systems.* MIT Laboratory for Computer Science technical report TR-389.

Tanner, J. M. and P. S.W. Davies (1985). "Clinical Longitudinal Standards for Height and Height Velocity for North American Children," *Journal of Pediatrics,* 107: 317-329.

Thissen, D. and R. D. Bock (1990). "Linear and Nonlinear Curve Fitting." In *Statistical Methods in Longitudinal Research, Volume II: Time Series and Categorical Longitudinal Data.* Academic Press.

Acknowledgments

Peter Szolovits has supplied many valuable comments on this work. Dr. Joseph Majzoub, Mario Stefanelli and Alberto Riva have also given helpful remarks. We thank Patient 002 and his family for the use of the growth data.