Collecting and Interpreting Temporal Data in an Expert System

William J. Long[1], Hamish Fraser[1, 2], Shapur Naimi[2], James Stahl[2]
[1]MIT Lab for Computer Science, Cambridge, MA,
[2]Tufts-New England Medical Center, Boston, MA

The collection and interpretation of temporal data in the cardiovascular domain is particularly important because the temporal characteristics of symptoms can provide important clues to the appropriate diagnosis. It is challenging because findings may be single events, continuous states, episodic events, or a progression described by more than one finding. This paper relates the nature of the problem and our experience over the past year as physicians have entered cases into the Heart Disease Program.

INTRODUCTION

Last year we reported on the Web interface for the Heart Disease Program[2] a program to diagnose the range of conditions that cause hemodynamic dysfunction[1,3]. Our experience over the past year as physicians have entered cases into the program has given us new insight into the problems of entering and interpreting the data. The input to the program is the same data that would be entered in the patient's chart including history, physical examination, and laboratory results. The information is entered into HTML forms using check boxes for categorical data and text boxes with bounds checking for numeric data (no free text). Because of the complexity of the data, the input is organized as a series of forms. The first form covers the basic information, from which a second form is generated for the needed details, and other forms as needed. The HDP then takes the input and generates a differential diagnosis using a pseudo-Bayesian inference mechanism enhanced with reasoning about the temporal and severity relationships involved.

The most challenging data, both for entry and for interpretation, is temporal information. For example, anginal chest pain is an important finding whose temporal pattern is crucial to determining an appropriate diagnosis. The causes of anginal pain are the various forms of cardiac ischemia: stable angina, unstable angina, and myocardial infarction (MI). The temporal characteristics of the pain are important for distinguishing the cause. Typically, stable angina presents with a chronic stable pattern of episodes of anginal chest pain lasting minutes brought on by exertion. A recent onset, worsening or change in the character of angina is indicative of unstable angina. Anginal pain longer and more severe, often still present when the patient is seen, is characteristic of an MI. However, not all pain with the typical characteristics of angina is ischemic in origin. Of course, the description given by the patient is likely to be somewhat different. The challenge for the HDP is to elicit from the physician a description of the pain that will allow the program to assign it with appropriate certainty to one or more of these causes.

If the physician specifies anginal chest pain on the initial input form, the next form asks the user to specify one or more types: exertional, at rest, or nocturnal. For each type specified, the user is asked the following details:

- TOTAL-DURATION-OF-SYMPTOMS: [select time from list]
- EPISODE-LENGTH: SECONDS, MINUTES, HOURS-OR-MORE
- FREQUENCY: ONCE, RARE, OCCASIONAL, OFTEN
- TEMPORAL-PATTERN: CHRONIC, NEW, WORSENING, IMPROVING
- CURRENTLY-IN-PAIN: NO, YES

The user can click an answer to each question. If WORSENING or IMPROVING are selected, the user then should select the time the change occurred.

With this menu, the user characterizes the chest pain and the program interprets the result as one or more appropriate findings. A (rare) example in which the user filled out all possible details is:

- anginal chest pain on exertion for 1 week, lasting minutes, occurring often, worsening for 1 week, but not currently in pain.
- anginal chest pain at rest for 2 days, lasting minutes, occurring occasionally, new onset, but not currently in pain.

The intent is that the total duration would indicate how long the patient has had episodes of anginal chest pain of the specified type (exertional, at rest, or nocturnal). Still, the user may be thinking in terms of the most recent exacerbation of the problem or on the other hand, in terms of all types of this symptom. Even in the example above, which seems to describe a patient who first had exertional angina a week ago progressing to angina at rest in the last two days, there is also a coronary angiography four months ago

with single vessel disease, a strong indication that the history of angina was much longer.

Often, the user will not fill in the details of the findings and the program is left with the task of making assumptions consistent with the rest of the case. Making assumptions is important because it is too time consuming to always fill in all details. This paper will illustrate the kinds of data entered by users and the kinds of reasoning required to use the data appropriately. First, we need to understand what information is provided by the temporal details.

INFORMATION IN TEMPORAL PROPERTIES

Several kinds of information are provided by the temporal properties of the findings. Each piece helps to constrain the diagnosis. The most obvious property is the temporal extent of the finding. If a finding has been present for a long time (months or years), the possible causes are diseases or pathophysiological states that can exist that long. A finding only present a short time does not provide as strong a constraint because all diseases start at some time, but if the cause is a disease with gradual onset, either the finding is mild or there is some exacerbating factor that worsens the finding and needs to be part of the hypothesis.

Often, past disease events are important in the cardiology domain. A previous MI often has long term effects on the heart. Aortic stenosis, even though corrected by a valve replacement, often leaves permanent ventricular hypertrophy. Even a disease like pneumonia increases the probability of a recurrence. Thus, many past events provide important context in which to interpret the current findings.

The second source of constraints is the order of the findings. Since effects can not precede causes, the findings related to the effects can not precede the causes. The reasoning can get complicated because not all findings are present the whole time the cause is present. Some findings take time to develop and some are only present over subintervals of the time the disease is present. This additional information also provides clues to the disease state. For example, if a finding requires time to develop, the cause must have been present for a sufficient time.

The pattern of the finding also provides important clues. The primary distinguishing patterns are acute onset (or single event), progressive, chronic stable, and episodic. Classic examples of these are: an MI with onset of symptoms so fast that the patient may drop in the midst of an activity; aortic stenosis which over the years progresses from a murmur to symptoms; exertional angina which is often predictable enough that the patient knows just how much exercise they can tolerate; and paroxysmal atrial fibrillation that occurs occasionally, usually unpredictably. Most diseases fall somewhere in between or have a range of possible behaviors. In addition, the patient description or the physician description may not precisely describe the appropriate pattern. Still, the observation of an acute, progressive, or gradual onset is useful evidence for diseases exhibiting those behaviors in the same way any other finding would be. For outpatients there may be few findings and temporal patterns may be the primary diagnostic clues.

Given that the HDP program needs extent, order, and pattern for the findings, the next problem is how to get them from the user.

TEMPORAL MENUS

Symptoms reported by the patient, medical history, medications, and often some test results carry temporal information needed for the HDP to generate consistent hypotheses. The chest pain details are an example of how this information is gathered by the HDP. The nature of forms for the Web interface requires the program to present the user with a number of questions at once. To limit the questions, each kind of data has a particular list of pertinent details.

Times are specified by selecting from a list: *now, 30min, 1hr, ..., 20yrs* with each approximately double the previous time. This is usually sufficient to specify the needed precision. The times are quantized to speed up the entry process and to let the user specify times at the minimum granularity needed. Often, the times are not known more precisely.

For tests and procedures HDP asks how long ago it was, since these are events without significant duration. This is a compromise because the time may not be precise enough to determine the order relative to other pertinent events. This issue commonly arises when there is a test and procedure in the past. Usually, the order is clear from the test result, assuming reasonable behavior of physicians at that time. Simple heuristics allow the HDP to determine the proper order. For example, an echocardiogram result not showing a mitral valve gradient must be after a mitral value replacement. The opposite situation is more problematic since a patient may be left with a gradient after surgery if the risk of further surgery is too great. Even so, assuming that the surgery corrected

the problem or that the user would tell you otherwise is a reasonable default.

For medications HDP asks only a duration. This ignores stopping, dosage changes, and other details. The essential information conveyed is the chronicity of the condition the medication is treating and the maximum extent of side effects and corrections caused by the medications. The duration is important because if the medication is effective, there may be few other findings indicating the presence of the treated condition. For example, a patient on diuretics for fluid retention will not have any evidence if the diuretics are effective.

For symptoms, extent, order, and pattern may all be important and each is different. Thus, the questions asked vary depending on the useful distinctions that might be conveyed. The questions for each symptom are a subset of the following:

- Total duration: [time list]: the time over which the patient has been experiencing these symptoms. This may be the time of the first event or the beginning of a continuous symptom.
- Event time: [time list]: for a symptom that occurs once, how long ago it took place. Diseases and symptoms that have taken place in the past can usually be viewed as events unless their effects require some temporal extent, e.g., corrected aortic stenosis existed for long enough to cause ventricular hypertrophy that remains.
- Episode length: seconds, minutes, hours-or-more; or minutes, hours, days-or-more (used for nausea): rough duration of the episodes. These are helpful for distinguishing causes.
- Last episode time: [time list]: this gives an indication of the acuteness of the problem. That is, whether it is in the past and can be viewed as a past event, recent and likely that the underlying cause is still present, or now.
- Frequency: once, rare, occasional, often: distinguishes between single and episodic events and gives a rough measure of severity.
- Change: stable, acute-onset, progressive: characterizes the episodic pattern. Progressive and acute onset are points on a continuum, so there is often uncertainty in the characterizations.
- Temporal pattern: chronic, new, worsening, improving: an alternate characterization of change, more appropriate for angina because it allows specification of the time of change.
- Worsening/Improving time: [time list]: the time of the change. For angina it is important to distinguish between changes within a couple of days and those over a week or two because the management of an acute MI, a recent MI, and unstable angina are different.
- Current state: present, absent: If the symptom is currently present, other symptoms need to be interpreted in that light. This is especially true of angina since the pain is not only an effect but a cause for other findings including the effects of sympathetic stimulation. Also, if anginal pain is present on examination, the episode is long enough to be unstable angina if not an MI.
- Type: various, depending on finding. The type helps to distinguish the causes, providing a finer grain of detail. This is important in cases with fewer symptoms.
- Associations: with chest-pain, nausea, diaphoresis, lightheadedness, etc. as appropriate for the symptom. This provides the important ordering information necessary to connect past symptoms and distinguish the diseases causing them.
- Severity: by whatever measure is appropriate for the symptom. Severity is often important for determining progression of disease as well as what other findings would be expected.

Each symptom has one to five of these questions so the user can describe the symptom and the HDP is better able to generate coherent causal hypotheses. These questions, as illustrated in the introduction, are very similar to the questions that would be asked of a medical student when presenting a case.

The menu is a compromise between the distinctions understood by the program and the language of the physician. Unfortunately, we have yet to find language to use in these menus that is unambiguous for all users that can handle the variety of ways physicians express temporal data.

The HDP takes this input and represents each finding as being true over an interval with possibly indefinite bounds and includes any properties provided by type and severity. Acute changes are handled by adding an additional finding with its own interval for the more severe period. Progressive findings must also be represented as multiple findings with different time periods because the change may be a new cause. That way, the progression can be attributed to the new cause and not the whole interval of the finding, which might precede the cause. For example, the introductory example produces the following:

- anginal chest pain on exertion [1 week to now], lasting minutes, etc.
- anginal chest pain at rest [2days to now], lasting minutes, etc.

- coronary angiography: single vessel disease [at 4 months]

The single vessel disease implies coronary artery disease, which the HDP assumes must have been present at least 2 years to be detectable. The week of exertional angina is most likely unstable angina and the angina at rest is either the unstable angina or possibly an MI.

CHALLENGES FOR INTERPRETATION

When the user completely specifies the findings for the HDP, there are still challenges for interpretation, an important one is deciding what findings are consistent with a diagnosis. For example, we originally ruled out ischemia as a cause for chest pain if the episode length was seconds instead of minutes. Unfortunately, we found that more than one user described pain that was clearly ischemic as lasting for seconds. Another question is how long after a change in anginal pain we can assume that this is a new pattern of stable angina rather than unstable angina. HDP originally restricted this to two weeks, but users have considered two months to still be an unstable pattern.

The pattern of the symptoms can often be difficult to interpret because progressive or acute worsening patterns can manifest themselves in several ways. The simplest is if the pattern is described in a symptom, e.g., progressive dyspnea on exertion. Often, the change is reflected as an addition or change of symptoms, e.g., the exertional dyspnea may become dyspnea at rest, or progress to exertional dyspnea with fatigue.

A common problem is that the user will enter both a mild and severe form of a symptom and only give the temporal extent of one of them. For example, he/she may specify angina on exertion for a year and angina at rest (without time). The HDP takes this to mean chronic angina on exertion and an acute change to angina at rest within the last day. This was originally implemented as both angina at rest in the last week and angina at rest within the last day to cover the important distinguishing time periods in which the change might have taken place, three different angina findings. The result was that the HDP would often diagnose both unstable angina and an acute MI in such situations, hypothesizing a progression for which the only evidence was the lack of time specification on the angina at rest. Assuming that the angina at rest took place within the last day still allows a diagnosis of unstable angina (at lower probability) so all possibilities are considered.

When only the severe form of the symptom is detailed, the mild form is considered to be chronic, that is, of more than a few months. This is a reasonably safe assumption because there are relatively few diseases that it eliminates.

Often the user leaves the details of the findings underspecified and the HDP must fill them in appropriately or occasionally modify details that are inconsistent. When this happens, the program adds a note to the input description printed with the differential diagnosis to alert the user. A typical example of the kinds of problems encountered is a patient with a mitral valve replacement 10 years ago, known diagnosis of mitral stenosis, and minor findings on echocardiogram but not including any mention of valve disease. Since physicians use the known diagnosis field to record any significant diseases that have taken place, often without specifying when they happened, this is reconciled by adding a time to the mitral stenosis so it is prior to the mitral valve replacement. In general, known diagnoses are taken to be historical (completed episodes) unless they are typically chronic diseases.

Sometimes the defaults are insufficient. In one case the user entered anginal chest pain (without time or type) and fatigue once 8 hours ago. The HDP assumes that unspecified anginal chest pain is chronic and therefore the program assumed that this was an incident of fatigue in someone who has coronary artery disease. Since fatigue by itself is common without significant disease, the HDP may leave it unexplained. However, the intention was to describe someone with new anginal pain and fatigue associated with it, a much more worrisome set of findings. This kind of problem illustrates why the HDP notes its assumptions when displaying the input.

Similarly, medications are often recorded without indicating how long they have been given. For medications typically given chronically, the HDP assumes that the medication has been given at least a week. This is an important assumption because medications are findings, especially for chronic diseases. This is appropriate for chronic diseases since the medications often cover the findings that would otherwise be present. It is not appropriate for diagnosing the acute changes in findings that did not influence what medications were chosen.

Often the chronic symptoms are not included in the description, as in the introductory example where there was probably chronic angina. This is handled by

putting a strong reliance on the test results and known diagnoses. Thus, a coronary angiography showing disease forces a conclusion of coronary artery disease. The program indicates to the user that angina was expected but not present, but uses the coronary artery disease as definite in putting together the rest of the hypothesis. This use of test results can cause problems if the users do not enter the results completely. Occasionally someone will enter an ejection fraction on an echocardiogram and as a result tell the program that the rest of the echocardiogram was normal. If the patient has a diastolic murmur, the HDP is at a loss to account for it.

DISCUSSION

The entry of patient descriptions is a challenging problem especially when the temporal characteristics of the findings are important for the program's reasoning. We reviewed a number of written case descriptions to see if there might be better ways of capturing the data. Unfortunately, it appears that the written descriptions are also full of ambiguities and often have less of the important detail than we have been able to collect from the physicians using the HDP interface. In those case descriptions adequately capture the temporal information, there is often a mixed strategy of describing findings and describing time periods. That is, there may be a short description of what happened a couple days ago, tying together several symptoms. This kind of strategy for gathering input data would require a much more flexible kind of interface.

The interface used in the HDP has proven useful for collecting an adequate description of the patient given the limitations imposed by the HTML forms. The questions must be asked in groups to limit the number of forms and it is unreasonable to require the user to fill in all of the possible details. Since details may be missing, it is important that the program be able to make appropriate about the acuteness of findings, the order of history findings and procedures, and the patterns they exhibit.

There is great variety in the cases the HDP must handle. For the patient in the hospital with complications of chronic diseases there are often many findings including laboratory results, detailed physical examination, and extensive history. In such cases, the challenge is to put together a consistent picture and to identify those parts that represent new problems. For the outpatient with a symptom or two, the only clues may be the patterns of the symptoms and the prior probabilities. The collection and interpretation of temporal data is critical to appropriate reasoning in both situations. A good interface has to be simple to use for the novice, encourage entry of the important data, and be time efficient (the physician's most valuable resource).

To overcome the current limitations, the interface will have to be more interactive, asking temporal questions sequentially and ensuring a consistent set of properties for the findings. In particular, the interface needs to ask about the order of findings with the same rough times. Also, the interface needs to identify and reask about important missing data, e.g., the anginal chest pain left unspecified or a murmur only identified as systolic. This directed interchange would be closer to what happens in the clinical setting where the intern is asked for pertinent details of findings and the relationships between them. There is a balance that will have to be maintained between forcing an order on the user and allowing the user to enter information in the most natural order. The more the interface meets the user's normal habits the higher quality the input is likely to be.

Acknowledgments

This research was supported by National Institutes of Health Grant No. R01 HL33041 from the National Heart, Lung, and Blood Institute.

References

1. Long WJ, Naimi S, Criscitiello MG. Evaluation of a New Method for Cardiovascular Reasoning. Journal of the American Medical Informatics Association. 1994;1:127-141.

2. W. J. Long, H. Fraser, and S. Naimi, Web Interface for the Heart Disease Program, In: Cimino, JJ, ed. Symposium on Computer Applications in Medical Care. Washington, DC, 1996 :762.

3. W. J. Long, H. Fraser, and S. Naimi, Reasoning Requirements for Diagnosis of Heart Disease, Artificial Intelligence in Med. (to appear).