Original Investigations

Research Paper ■

Evaluation of a New Method for Cardiovascular Reasoning

WILLIAM J. LONG, PHD, SHAPUR NAIMI, MD, M. G. CRISCIANELLO, MD

Abstract **Objective:** Evaluate the accuracy of the detailed diagnostic reasoning of the Heart Failure Program incorporating a new mechanism to handle temporal relationships and severity constraints.

Design: Tools were developed to summarize diagnoses and automatically generate evaluation forms. Five expert cardiologists were asked to review the reasoning of the program, with two analyzing each case. Cases were gathered retrospectively for diversity and difficulty and 26 randomly selected cases were evaluated. The underlying issues were identified and classified.

Results: Both reviewers rated the first diagnosis correct in 25% of the cases and at least one rated it wrong in 10%. Analyzing the detailed reasoning, 137 issues were raised, about 5.3 per case. Of these, 53% were possible concerns raised by one reviewer. Of the 5.3 issues per case, 2.5 were attributable to controversies, misunderstandings, or mistakes; 1 was due to the overly simplistic representation of the summaries; and 1.8 were issues related to the program.

Conclusion: Overall, the program is capable of providing high-quality detailed diagnostic hypotheses for complex cardiovascular cases. The results highlight several issues: 1) the difficulty of effectively summarizing hypotheses, 2) the nature of a physician's causal explanation, and 3) some problems in evaluating detailed diagnostic reasoning. The mistakes the program made imply that some additional refinement is needed but that the reasoning mechanisms developed can support the appropriate reasoning. The appropriate next step is a prospective evaluation addressing the program's usefulness.

■ *J Am Med Informatics Assoc.* 1994;1:127-141.

The Heart Failure Program (HFP) is a computer program designed to assist physicians in reasoning about patients with cardiovascular disease.¹⁻³ It takes as input the description of the patient's history, physical examination findings, and test results at the level of

detail one would find in the physician's description in the patient record. This information is used by the program to generate a differential diagnosis for the case. It attempts to generate a hypothesis corresponding to each likely combination of diseases that could account for the findings. The program uses a pseudo-Bayesian network as a knowledge base representing the physiologic causality of the cardiovascular hemodynamics from which it generates hypotheses that explain the findings. A small portion of the network is sketched in Figure 1. Each arrow between nodes represents a causal link with a conditional probability. Each hypothesis is a subnetwork of the knowledge base instantiated with the findings of the case. This subnetwork is a complete explanation from primary causes through the pathophysiologic mechanisms to the findings. The correspond-

Affiliations of the authors: MIT Laboratory for Computer Science, Cambridge, MA (WJL); New England Medical Center and Tufts University, School of Medicine, Boston, MA (SN, MGC).

Supported by National Institutes of Health Grant No. R01 HL33041 from the National Heart, Lung, and Blood Institute and Grant No. R01 LM04493 from the National Library of Medicine.

Correspondence and reprints: William J. Long, PhD, MIT Laboratory for Computer Science, 545 Technology Square, Room 420A, Cambridge, MA 02139.

Received for publication: 7/6/93; accepted for publication: 11/10/93.

Figure 1 A portion of the Heart Failure Program knowledge base.

ing portion of a hypothesis is shown on the left side of Figure 2. The only findings not in the hypothesis causal network are those considered normal or better explained by something outside the domain. For each hypothesis generated, the program can compute a probability by combining the probabilities in the network. The differential diagnosis consists of these hypotheses ordered by the computed probability. To limit the differential to likely hypotheses, the list is cut off when the probabilities fall below 1% of the best hypothesis. Thus, the differential may consist of one or many hypotheses. This type of diagnostic hypothesis is much more informative than those of earlier programs, such as AI/RHEUM,⁴ QMR,⁵ or DXplain,⁶ which provide hypotheses consisting of single words or phrases. The structure of the HFP hypothesis explains the findings by showing the causal mechanisms producing them and thus provides a justification that the physician can evaluate to decide whether the conclusions are reasonable. The causal mechanisms are modeled at a level of detail consistent with that at which a physician might explain the findings to a colleague. This enhances the usefulness of the hypotheses for justification because the concepts fit naturally with the understanding of the physician. It is not a handicap for reasoning because this is presumably the level at which the human physician does similar reasoning and the physician is the best model we have for this kind of reasoning.

We previously conducted a formative evaluation of the program.⁷ This evaluation used 242 cases collected from discharge summaries and compared the program diagnoses with diagnoses collected from cardiologists using the same information. We examined the maximum potential accuracy of the program by iteratively revising the knowledge base and rerunning the cases. Ultimately, the program was able to diagnose 90% of the cases correctly. The main reasons that the remaining 10% could not be diagnosed correctly were the lack of reasoning about temporal relationships, the distinctions of different severities of diseases, and inappropriate combining of probabilities of multiple diseases. As a result of this evaluation we have developed methods for reasoning with the additional constraints provided by temporal and severity relationships to deal with some of the limitations encountered. We have also developed an appropriate framework for computing the probability of multiple diseases.⁸ This paper presents an overview of the changes that have been made to HPF and reports an evaluation of HPF's performance.

Temporal Relationships and Severity Constraints

The additions to the program include both enhancements of the knowledge base and changes in the reasoning methods. In the knowledge base the dis-

Figure 2 Summarization of part of a hypothesis.

eases and pathophysiologic states (all referred to as nodes) are subdivided into levels of severity and subtypes with additional constraints on the causal links. The levels of severity and subtypes distinguish qualitatively different hemodynamic consequences. The link constraints reflect what severities and time bounds are required for the cause to produce the effect. For example, aortic regurgitation has three levels of severity reflecting situations in which there is a murmur without any hemodynamic consequence, situations in which the heart has compensated by dilating, and those in which the aortic regurgitation has caused deterioration of systolic function. These differences are reflected in the different probabilities and constraints on the links conditioned on the severity. Aortic regurgitation is also divided into two subtypes: primary valvular regurgitation and secondary regurgitation due to dilation of the aortic root. This distinction allows the effects to be different or have different probabilities for different subtypes. General temporal constraints on the nodes reflect limitations on causality, such as how long it takes for the nodes to become true [e.g., minutes for an acute myocardial infarction (MI) to years for aortic stenosis], how long the node will remain true after the cause ends, and how long a patient might have the state (e.g., an "acute" MI only exists for two days by definition before it becomes a "recent" infarct). The causal links include statements relating the severities and temporal constraints. For example:

```
(defnode (high pulmonary_vascular_resistance)
  severities ((1+ reversible)(2+ irreversible))
  causes ( . . .
    (pulmonary_embolism
      :prob (3+ 1.0 → 1+ .3 → 2+))
    ((low arterial_po2) :prob ((< 1wk) .3 → 1+))
    ((high la_press)
      :prob (3+ (> 1yr) .2 → 1+ .2 → 2+))))
```

Thus, high pulmonary vascular resistance (PVR) has two severities, distinguished by being reversible and irreversible, rather than a particular measured resistance. A number of possible criteria could be used for defining severity, such as right ventricular dilatation, but reversibility was chosen because in the case of PVR it is correlated with severity and has important implications for management. The possible causes include pulmonary embolism, which if severe (3+) causes high PVR immediately and is irreversible 30% of the time. Low arterial oxygen levels for a limited time (a week or less) can cause reversible high PVR 30% of the time. Finally, chronic (a year or more) severe (3+) high left atrial pressure (as mitral stenosis might cause) can cause either severity of high

PVR. The severities of pulmonary embolism and high left atrial pressure are also defined in terms of qualitative differences in their effects.

These constraints are enforced by the reasoning mechanisms at two levels: 1) The causal pathways that are computed when the knowledge base is loaded are pruned using the constraints, and 2) the patient-specific nodes that are generated from the findings to build the hypotheses carry the constraints that guide the process of building the hypotheses. In the completed hypothesis, each node has one or more data structures called *severities* that represent a constraint on disease severity and subtype true over a specified temporal interval. For example, a *severity* for high PVR includes the following:

```
[severity node: (high PVR) min: 1+ max: 1+
 type: none
 justification: path to supporting findings
 time: time interval with ranges for beginning and ending
 causes: list of causing severities
 effects: list of severities this causes]
```

Each of these structures is linked to the corresponding *severity* structures of the causes and effects. More than one *severity* structure is necessary when the evidence for a node supports different constraints on severity over different time intervals, such as a chronic disease with acute worsening.

Methodology

A number of evaluations of medical expert systems have been reported, mostly focusing on the accuracy of diagnostic information.^{4, 5, 6, 9, 10} Evaluating the accuracy the program is the first step toward establishing the clinical usefulness of the program. The essential difference between these evaluations and the evaluation needed to assess the accuracy of the HFP stems from the difference between the diagnoses—a fixed list of atomic diagnoses in the case of the other diagnostic programs, versus a network of causal relations for each diagnostic hypothesis for the HFP. To some extent, this more extensive set of possible conclusions was also faced in the evaluation of therapeutic decision support systems.^{11, 12} However, here too there was a single decision for the evaluators to judge. The hypothesis structure presents a much more difficult challenge for evaluation. With atomic diagnoses it is relatively clear what is correct or incorrect and it is easy to compare expert diagnoses with the program's diagnosis. For the HFP we needed to break the hypotheses down into manageable units and develop a method for evaluating the hypotheses.

While we were developing the temporal and severity reasoning and knowledge base, we used a learning set of 92 cases for testing and refining the algorithms. These were used to ensure that the additions to the knowledge base were consistent and captured the distinctions pertinent to the domain.

To evaluate the program, 60 cases not previously seen by the program were gathered independently. The cases were selected to have as much variety as possible to present the program with a broad challenge. The cases included patients under care at the time of collection and as well as records of previously treated patients not used for program development. To evaluate the program at the time point likely to be the most challenging for diagnosis, each case was entered using data available at a time after the initial physical examination was completed but before new echocardiographic or catheterization data were available. Any echocardiographic or catheterization data obtained earlier were entered, because the admitting physician would have had them available. Thus, the evaluation focused on the more difficult issues of reasoning from physical findings. Once the cases were collected and entered, the files were set aside until we were ready to conduct the evaluation.

We developed a format for evaluating the cases by using the learning set. Because each complete hypothesis consists of 20–40 nodes plus the findings and each case has one to several hypotheses in the differential, it was necessary to distill the information into a format that would be more manageable for the reviewers. To accomplish this we developed an algorithm for automatically summarizing the hypotheses. Summarization is possible because the nodes in the hypothesis play different roles. While some characterize the state of the patient, others give details of the causal mechanisms. For justification and explanation the mechanism nodes are useful, but when conveying the essence of the diagnosis is important, they can obscure the more important nodes and increase the size and apparent complexity of the hypothesis.

The summarization algorithm uses node attributes (recently added to the knowledge base) that identify their roles relative to other nodes. Summarization retracts the mechanisms into the nodes that classify the process. That is, the classifying node covers any nodes that represent mechanisms of it. Part of the summarization of the case in the appendix is illustrated in Figure 2. The node labeled *regional flow deficit* has the attribute *mechanism* so it is clustered under its cause (*fixed coronary obstruction*). The myocardial infarction node has an attribute indicating that its

cause, myocardial ischemia, is a mechanism. Next the definitions of *summary concepts* are checked to see if there is any further clustering to do in the hypothesis. Here, *coronary artery disease* is defined to cover *fixed coronary obstruction*, *old MI*, and *exertional angina*. Multiple applications of the summary concepts are possible. For example, if *low cardiac output* and *pulmonary congestion* are present, they are clustered together as *left heart failure*, and if both *left heart failure* and *right heart failure* are present, they are clustered together as *congestive failure*, as in the summary in the appendix. Once the nodes are clustered, they are given the name of the summary concept or the primary node in the cluster. The causal links between the clusters are formed from the links that existed between nodes in the different clusters. Some heuristics are invoked to link the clusters appropriately. For example, since *nitroglycerin* is an effect prior to the primary node in the MI cluster, it is linked to the *coronary artery disease* cluster. This summarization mechanism greatly reduces the size of the hypotheses.

The alternative hypotheses are just like the primary hypothesis, but with lower computed probability. Thus, they are summarized in the same way. However, because they usually contain many of the same nodes as the primary hypothesis, differing in only a few places, it is more compact to describe them by telling how they differ from the primary hypothesis. This is described in terms of the differences between the summaries rather than the differences between the underlying nodes to maintain a consistent vocabulary for the evaluators. For the evaluation, we took one further step and merged the difference lists of the alternative hypotheses into a single list. This step may lose some information about what combination of additions and removals go together, but still gives the evaluator the range of alternatives suggested by the program.

The forms given to the reviewers were automatically generated by the HFP using the summarization mechanism (see appendix). These forms were organized as follows: The first page of the form had only the input description given to the program, allowing the reviewer to consider the findings before seeing the program's conclusions. The second page had the graphic display of the summary of the best hypothesis of the program followed by a textual description of the summary, outlining the causal relationships, so the reviewer could use whichever form was easier. Whenever there were alternative hypotheses generated by the program, the features of these were summarized as a list of what they added, changed, or removed from the best hypothesis.

Starting with the next page, the reviewers were asked to judge the results. The summaries of the best hypothesis and of the alternative hypotheses were melded together and then divided into a statement about each node in the summary. These statements included what was coming in to the node (causes or worsening factors) and what was going out (findings and other nodes caused, possibly caused, or worsened). Because these statements cover multiple hypotheses, they include phrases such as "possibly caused by." The alternative hypotheses were combined with the primary hypothesis in this way to give the reviewer the range of possible causes and effects for each node. This gave a more detailed examination of the alternatives than reviewing the difference list. After each statement the reviewer had five choices with the following meanings:

Correct	This is a reasonable accounting of the findings and nodes.
Possible	This is a possible interpretation of the findings and nodes, but not the best.
Partly correct	This is mostly reasonable, but has some small subset of the findings or nodes wrongly attributed.
Wrong	This node in the hypothesis is wrong.
Seriously wrong	This node in the hypotheses is wrong and could lead to inappropriate care for the patient.

For those statements that the reviewers determined to be less than correct, they were asked either to circle the offending part or to make a short comment to help us analyze the problem.

Other decision classifications have been used in evaluations, such as the *ideal*, *acceptable*, *suboptimal*, *unacceptable* scale used by Hickam et al.¹² Our scale is similar, but does not try to differentiate between acceptable and suboptimal. Instead, we attempt to capture the difference between a statement that is suboptimal because it has some part (a cause or effect) that the reviewer disagrees with (*partly correct*) versus the statement itself not being the best explanation (*possible*). The difference between *wrong* and *seriously wrong* was intended to highlight errors that could have consequences if acted on, because the statements vary widely in importance.

The first statements to be reviewed were those for which the summary node was in all hypotheses. These were followed by statements that were in only some of the hypotheses and thus represented aspects of the alternatives. Following these statements was space

to indicate any diagnostic or pathophysiologic determinations that the reviewers felt had been left out. Finally, the textual statement of the best hypothesis and the summary of alternatives were presented with the same choices. These provided the overall assessment of the diagnostic performance of the program on the case.

Because we recognized some problems with the summaries, to help resolve questions we also included as a last page the detailed graphic display of the leading hypothesis (not included in the appendix), but such hypotheses were consulted only in rare instances.

The reviewers were five senior cardiologists with no connection or previous exposure to the HFP. All are in active practice at teaching hospitals, and they represent the four medical schools in Massachusetts. A few days before they were given the study cases, they received instructions and a sample evaluation form from the learning set to familiarize them with the task. The evaluation was conducted as a single three-hour session with all reviewers together, but working individually. That format provided us a way of answering the questions they had while critiquing the cases.

The cases were arranged in random order and given one at a time to the cardiologists. Each case was reviewed independently by two of the five cardiologists. Each cardiologist reviewed as many cases as possible in the three hours. The numbers completed by the individual cardiologists varied widely, ranging from seven to 17. The second reviews of the first ten cases were held until the second hour to balance out any differences due to initial adjustment to the format. The numbers of statements in individual cases reviewed by the cardiologists, including the statements about the leading hypotheses and about alternative diagnoses, ranged from six to 15 and averaged about 11. The 26 cases that were completed with two reviews each were used for subsequent evaluation.

The 26 cases reviewed included a wide variety of diseases, and most of the patients had multiple diseases. There were five acute MIs, two recent MIs, two cases of unstable angina, and nine others with stable coronary artery disease. Valvular disorders included nine with mitral regurgitation; four each with aortic stenosis, mitral stenosis, and tricuspid regurgitation; three with aortic regurgitation; and one with aortic and mitral valve replacements. Five patients had cardiomyopathy. Two had significant hypertension and two had pulmonary hypertension. There were also cases of endocarditis, pericarditis, and pos-

Table 1 ■

Reviewer Pairs' Overall Judgments of the First Hypothesis and Alternatives

Judgment of Reviewer 1	Judgment of Reviewer 2	Hypothesis 1	Alternatives
Correct	Correct	7	3
Correct	Possible/partly correct	4	13
Possible/partly correct	Possible/partly correct	12	5
Seriously wrong	Possible/partly correct	2	0
Wrong	Wrong	1	0

sible aortic dissection. Renal insufficiency and pulmonary disease were well represented. The first hypotheses for the cases had an average of 3.2 disease nodes. These were all complex cases, chosen to be difficult, but otherwise representative of the case load of a tertiary care hospital.

The reviewers were not asked to record or critique their own diagnoses. Having the reviewer's diagnosis evaluated as well, as has been done in earlier evaluations,¹¹ would make it possible to rank the program relative to the evaluators. However, the additional work required of the reviewers would have greatly increased the time required per case, but more importantly, the nature of the human differential diagnoses is sufficiently different from that of the computer-generated diagnoses to make it very difficult to compare them. Physicians vary the specificity of the diagnosis depending on factors such as certainty and they tend to focus on the acute problem, sometimes not mentioning chronic problems. Additionally, because the source of a diagnosis is easily recognized, the bias of knowing whether the reviewer is critiquing a physician or the computer would further obscure the comparison. As a result, how the reviewers would judge each other will remain unknown, although we can infer from the explicit disagreements in evaluating the cases that there would be significant differences among the reviewers.

Because the evaluation was held as a group session, we were able to answer many of the reviewers' questions and avoid a number of misunderstandings. They were also encouraged to write comments and circle inappropriate parts of statements. As a result, it is possible to analyze the evaluation to determine the reasons for dissatisfaction with statements they found less than correct. The environment also generated some conversation that disclosed that some concepts that we previously had thought were widely accepted were controversial.

Results

The statements that were judged can be divided into those about the overall diagnosis and those about the details. Overall rating was done on the first hypothesis in each case and on the summary of the alternatives in the 21 cases that had additional hypotheses. The ratings are given in Table 1. There are three cases in which at least one of the two reviewers who did the assessment considered the first hypothesis wrong. In two of these, the program missed diagnoses of mitral stenosis, once because the location of the murmur matched tricuspid stenosis only and once because the program considered the murmur to be a flow murmur of the mitral regurgitation that was present. In the third case, the program left pleuritic chest pain unaccounted for. The reviewer criticized it for not suggesting pulmonary embolism, although it was actually pericarditis. The two *seriously wrong* judgments were made by different reviewers. The judgments of the alternatives included fewer in complete agreement and none considered wrong. With only about 10% of first hypotheses judged by any reviewer to be wrong, there seems to have been general agreement, but with only 25% of first hypotheses judged by both members of the reviewer pair to be completely correct, and fewer of the alternatives, there was considerable disagreement about the details.

In this analysis and in the analysis of the detailed statements, it became clear that some of the distinctions in the rating scale had no practical significance. The reviewers seemed to use *possible* and *partly correct* interchangeably, so in the following analysis they are combined and called *possible*. There was no statement that both reviewers in a pair rated *seriously wrong*; in fact, there was no statement that one reviewer rated *seriously wrong* that the other reviewer rated even *wrong*, so *wrong* and *seriously wrong* are combined as *wrong*.

To determine whether there was any systematic bias among the reviewers, we compared the fractions of the statements each rated as correct, possible, and wrong with the fractions with those ratings as judged by the rest of the reviewers. We also compared the ratings of statements about disease nodes in all hypotheses versus those in some hypotheses. Those fractions with their levels of significance and the fractions of the total statements they represent are given in Table 2. This analysis indicates that the reviewers were well balanced. The two reviewers who differed significantly from the rest contributed the smallest numbers of ratings and their differences balanced each other. Comparing definite nodes with possible nodes, the possible nodes were judged more harshly.

Table 2 ■

Distributions of Ratings by Reviewer and by Type of Statement*

	Reviewer					Type of Statement	
	1	2	3	4	5	Definite	Possible
% correct	48	42	66 (0.001)	46	30 (0.01)	54 (0.001)	37 (0.001)
% possible	40	42	22 (0.001)	47	54 (0.05)	38 (0.05)	48 (0.05)
% wrong	12	15	11	7 (0.05)	15	8 (0.01)	16 (0.01)
% TOTAL	16	23	15	34	12	45	29

*Statements true of all hypotheses are definite, those true of some hypotheses are possible.

Analysis of the detailed statements requires caution because each statement may overlap with one to several other statements in the same case. In relating a node to its definite and possible causes and effects, each statement touches on several nodes. If there is a node that the reviewer feels should not be in the hypothesis, that affects his judgment of statements about each of the causes and effects as well as the statement about the node itself. For that reason we analyzed the critiques to determine the sources of negative judgments and clustered the affected statements together. This process required some interpretation in many cases, but fortunately the reviewers' comments were extensive enough that it was always possible to have a good idea what they were concerned about.

In the 285 detailed questions there were 137 issues raised by one or both of the members of the reviewer pairs. One might infer that the two reviewers of the pair were in complete agreement with 148 of the statements, but because of the influence of the issues on the statements about causes and effects, there were only 92 statements that were considered correct by both reviewers. Each issue was assigned for each reviewer the most serious rating of any of the statements in which that issue arose. Some issues were comments from the reviewers about relations and hypotheses that they considered missing. Because these did not have ratings, they were all rated *possible*. Thus, if *possible* = P and *wrong* = W, the issues that concerned both reviewers could be rated WW, WP, or PP, and those that were of concern to only one would be rated W or P. The 137 issues that arose in the cases were rated as follows: WW 11, WP 10, PP 16, W 27, P 73. The majority (53%) were possible changes that were of concern to only one reviewer. Still, there were 11 issues that the reviewers agreed were wrong and another ten that one considered wrong and the other reviewer thought were not the best choices.

To determine the source of these issues, we analyzed and classified them. Classifying the issues is open to

some interpretation, but it is very useful in determining whether the issues imply that refinement of the program is needed, the method of summarization and presentation misrepresented the conclusions of the program, there was some misunderstanding or mistake on the part of the reviewer, or there is an underlying difference of opinion among cardiologists. To avoid missing problems with the program, we classified issues as relating to the program whenever there was any doubt.

The classifications we used are as follows:

Controversy. These are issues in which there was clearly a difference of opinion among cardiologists. Several of these came to light during the evaluation in discussions in which the reviewers disagreed with each other. Others are disagreements with carefully considered representations in the program, that is, disagreements between a reviewer and the developers. These were reviewed and classified by the cardiologist developers.

Reviewer wrong or inconsistent. These are issues where the criticism is in conflict with the actual patient state as indicated in the patient record or, in two cases, the reviewers' ratings of different statements in the same case are mutually inconsistent. Given the limited information in the input, several diseases may be appropriate to consider in addition to what the patient actually had. Therefore, the judgment was classified as wrong only when the program's statement corresponded to the actual situation and reviewer rated that as incorrect.

Misunderstanding. These are issues in which the reviewer probably overlooked part of the information or misunderstood some of the (occasionally convoluted) automatically generated text. That is, the intended meaning of the statement is consistent with the reviewer's objection.

Summarization. These are issues in which the program that summarized the hypotheses and put the information from multiple hypotheses together into single

Table 3 ■

Classification of the Issues Raised by the Reviewers

	Reviewer Pair's Judgments					Total
	Wrong/ Wrong	Wrong/ Possible	Possible/ Possible	One Wrong	One Possible	
Controversy	1	4	3	6	27	41
Reviewer wrong	0	0	0	2	4	6
Misunderstanding	0	0	1	5	10	16
Summarization	2	0	3	6	15	26
Program	8	6	9	8	17	48
TOTAL	11	10	16	27	73	137

statements about nodes obscured the relationships that exist among the nodes or inappropriately labeled node clusters.

Program. Relationships in the program that need to be reexamined. (These are further classified below.)

As an example of how the cases were analyzed, consider the evaluation form in the appendix. The first reviewer marked all statements *correct* except two: one about atrial fibrillation and one about atrial septal defect (ASD), which were marked *possible*. The second reviewer marked the atrial fibrillation and mitral regurgitation statements *partly correct* with a note that atrial fibrillation is caused by dilated cardiomyopathy, not mitral regurgitation. He marked the statement about dopamine *wrong* and those about acute MI, coronary artery disease, and ASD *possible*. Thus, these reviews resulted in one W: dopamine contributing to heart failure; two PPs: atrial fibrillation caused by dilated cardiomyopathy instead of mitral regurgitation and ASD causing fixed splitting S₂; and one P: acute MI as evidenced by elevated CPK-MB (with the coronary artery disease statement considered to be the same issue). These were classified as follows: ASD, because it is a weak alternative hypothesis, was attributed to the program; the dopamine and the cause of atrial fibrillation were considered controversies, because the cardiologist developers consider them correct; and the acute MI was a misunderstanding. The acute MI is an example of the kind of misunderstanding that can arise from an automatically generated evaluation form. The acute MI was included in only the first of the two hypotheses and therefore listed with the nodes present in some hypotheses only. The reviewer considered acute MI a likely possibility, and therefore was in agreement with the program, but marked the statement *possible* rather than *correct*, confusing the correctness of the statement with the likelihood of acute MI.

Given this scheme, the 137 issues identified by the reviewers were classified as shown in Table 3. One

controversy that accounted for six of the issues was whether diastolic dysfunction causes low cardiac output and left heart failure, that is, whether a patient with LV hypertrophy, a normal ejection fraction, fatigue, and pulmonary congestion should be described as having diastolic dysfunction causing the findings. It is clear that patients, especially older patients, present with those findings, but it depends on how one defines "diastolic dysfunction" whether that is the cause or not. This difference generated discussion among the reviewers. Several other controversies were probably issues of definition as well—whether mitral stenosis causes anginal chest pain, what constitutes a left atrial abnormality on electrocardiogram, or how broadly one may define chronic obstructive pulmonary disease (COPD). Some of the disagreements were about whether a particular disease was adequately supported, for example, whether murmurs as the only direct evidence were sufficient to suggest aortic stenosis or tricuspid regurgitation. Others were findings that the program left unaccounted for that the reviewers felt should be accounted for, such as cough and nonspecific ST- and T-wave changes.

Assuming that different reviewers take different sides in controversies, make different mistakes, and have different misunderstandings, the first three categories in Table 3 are issues that would have a high rate of disagreement between reviewers. Because these are only a sample of what might contribute to inter-reviewer disagreement and they are 45% of the issues, the differences among the reviewers are likely to be comparable to the differences the reviewers had with the program.

Most of the summarization problems resulted from using the phrases *caused by* and *accounted for* in the statements to represent all of the linkages between summary nodes. Once clusters of nodes are abstracted to summary nodes, many of these become influences not normally considered causality (e.g.,

the summary saying that aortic stenosis is causing hypertension because the hypertension cluster includes the high LV pressure node which is also caused by aortic stenosis) or clusters having inappropriate labels given the severity or what they were influencing (e.g., a cluster labeled high blood volume causing elevated liver function tests, skipping the intermediate splanchnic congestion). Another summarization problem was grouping acute and chronic manifestations of nodes together. For example, the program hypothesized that one patient had chronic mitral regurgitation worsened by an acute MI. In summarizing the mitral regurgitation it listed the acute MI as a cause and included both acute and chronic findings as effects. Another problem was not having different names for different severities of the clusters. For example, situations ranging from mild tachypnea to frank pulmonary edema were all labeled *left heart failure*, whereas treated failure without symptoms or with minor symptoms should be called *compensated left failure*.

Criticisms Related to the Program

The remaining 35% of the issues were related to the functioning of the program. These can be further classified as shown in Table 4. The largest source of these stemmed from a difference between the natures of hypotheses from a Bayesian network and hypotheses that a human expert would generate. The single criterion for a good hypothesis from a Bayesian network is high probability. Our heuristic approach also tends to find single causal pathways unless there is separate evidence for more than one. The criteria implicitly used by experts in generating hypotheses require adequate evidence for diseases, inclusion of all important diagnoses, and inclusion of all plausible explanations for findings. They handle side effects of therapies in a special way, liberally considering the possibility but being conservative about definite attribution. These differences come in conflict in a

number of ways. First, the program does not require any particular evidence to invoke a disease. In several instances the program included diseases with inadequate evidence (from the reviewer's perspective) because that yielded a higher-probability hypothesis, e.g., concluding low renal perfusion on the basis of the renal function tests without evidence of pedal edema, suggesting anemia from findings without having a hematocrit, or using functional tricuspid regurgitation to account for a murmur when there is cardiomegaly. Second, the program will eliminate important diseases if the probability is too low. The program left out pulmonary embolism because the pleuritic chest pain has a high false-positive rate and the shortness of breath could be otherwise explained, not taking into account the importance of diagnosing pulmonary embolism, which is potentially serious and very treatable. Third, the program also left out possible causal pathways for findings such as cough, LVH on electrocardiogram, and pulmonary hypertension because one reasonable explanation had been found. In one case, poor ventricular function was attributed to existing hypertensive heart disease and not to an acute MI because all of the findings were consistent with chronic disease. The reviewers disagreed because the acute MI will certainly worsen the ventricular function even if it did not cause it and even though there is no actual evidence of acute worsening. Finally, in a few instances the program attributed findings to drug side effects when there were other possibilities, e.g., steroids to account for pedal edema when there were possible cardiac causes.

These differences do not imply that the Bayesian network is incompatible with expert-like hypotheses, but it does imply that some additional processing is needed. First, a utility model is needed to avoid pruning out important hypotheses with probabilities below the threshold. This could be implemented by extending the reasoning mechanism to include value nodes as in an influence diagram or by doing some

Table 4 ■

Classification of the Criticisms Caused by the Program

	Reviewer Pair's Judgments					Total
	Wrong/ Wrong	Wrong/ Possible	Possible/ Possible	One Wrong	One Possible	
Needs more evidence	2	1	4	3	2	12
Other causes for finding	2	1	1	0	4	8
Missing important disease	1	0	2	0	0	3
Missing concept	0	0	1	0	5	6
Severity definition	0	2	0	4	2	8
Suspect probabilities	3	2	1	1	4	11
TOTAL (program)	8	6	9	8	17	48

post-processing of the hypotheses to allow the appropriate reasoning about utilities. Second, additional pathways to findings need to be added to reflect alternative causes within the hypothesis. Third, diseases that are not well justified usually represent the most likely of several possible causes for a set of findings. These should be presented to the user in a more general form, using names such as "poor ventricular function" that leave the actual causes unspecified. Finally, therapy side effects should be handled in a way that considers the additional means the physician has of determining the causes by changing or adjusting the therapy.

A few of the alternative diagnoses suggested by the reviewers involved concepts that we have not included in the knowledge base, although we have considered adding them. These were thyroid therapy causing hyperthyroidism, high PCO_2 from inappropriate ventilator setting, and left ventricular aneurysm. These are concepts that can be added to the knowledge base. There were other suggested alternatives that still seem inappropriate, such as using diabetes as an explicit cause (beyond changing some of the probabilities, which it currently does) and cerebral embolism.

Another source of criticisms was inappropriate severity constraints. These included the level of hematocrit needed to account for high cardiac output, the severity of renal insufficiency that causes nausea, and the degree of pulmonary hypertension that causes RV hypertrophy. These can all be corrected in the knowledge base.

The rest of the issues are attributable to the probabilities in the knowledge base. There were diseases or pathophysiologic states in hypotheses deemed too unlikely to be considered (but not ruled out) and diseases and states that the reviewers felt should have been included. One problem that occurred twice was determining the source of murmurs from their location. The worst example was concluding the rare disease, isolated tricuspid stenosis, because the murmur description did not fit anything else. We have done considerable work on murmur description, but the variability of descriptions encountered in cases continues to be a challenge. One solution is to give the locations a high error rate. In two situations unlikely causal mechanisms were proposed by the program. One was essentially overcompensation: high vagal stimulation (evidenced by diaphoresis) causing low vascular resistance (with low blood pressure) causing high cardiac output (with a flow murmur). Another was letting the wrong causal pathway dominate in determining the effects of septic shock. The

program concluded that the decreased filling pressure would cause low cardiac output rather than the decreased systemic resistance causing increased (or normal) cardiac output. Both of these problems can be corrected by adding causal constraints. Still, they provide interesting examples of how the underlying assumption of conditional independence in Bayesian networks can misrepresent causal models if secondary dependencies are overlooked.

One very general problem is the use of pertinent negatives to justify conclusions. The program uses negatives to find appropriate causal pathways in building hypotheses and to determine their probabilities, but only the abnormal findings are included on the displays and used to justify the nodes. It was clear from both written and oral comments that in many cases the negatives were crucial to making the statements convincing. In one case, the program had coronary spasm as the primary cause for an MI because a catheterization done six months prior had not mentioned the coronaries, which the program interpreted to mean that they were clear. Thus, coronary spasm was the most likely cause. Without that assumption, coronary artery disease would be a more likely cause. In another case, one reviewer judged the lack of aortic stenosis in the hypothesis to be a serious error and the other reviewer judged the same statement to be correct and commented that it should be justified by the lack of LVH on the electrocardiogram.

Thus, we have determined that the first hypotheses were rated correct by both reviewers in the pair in 25% of the cases and wrong by at least one in 10%. Analysis of the detailed reasoning identified 137 issues, about 5.3 per case. Of these, 53% were possible concerns raised by one reviewer. Of the 5.3 issues per case, 2.5 were attributable to controversies, misunderstandings, or mistakes; 1 was due to the overly simplistic representation of the summaries; and 1.8 were issues related to the program. All of the program issues can be handled with refinement of the knowledge base and some additional processing to incorporate concerns of utility and evidence.

Discussion

This evaluation provides us with a reasonable picture of the diagnostic performance of the Heart Failure Program. The number of cases in the evaluation was too small to measure the performance relative to physicians. In any case, that would be missing the point. The performance of the program is at a level such that the experts can relate to the explanations and their criticisms are mostly in the details of the di-

agnoses. There are still occasional serious errors made by the program, but they imply that further refinement of the knowledge base is needed rather than fundamental changes to the reasoning mechanisms. The severity and temporal constraints give the program the tools necessary to keep from generating impossible causal pathways in the hypotheses and allowing better discrimination of the likelihood of the pathways it does generate. Some additional reasoning is needed to create a better fit with the human expectations of good hypotheses.

The evaluation also illuminated a number of difficulties faced in evaluating a program that provides highly detailed assessments. The first issue the evaluators commented on was the artificiality of the case descriptions with most of the echocardiographic findings left out. We chose to rely mostly on physical examination, electrocardiographic, and X-ray findings to get a more extensive test of the program's reasoning, but cardiology in the United States has come to a point where in complicated cases cardiologists are reluctant to form conclusions without seeing the echocardiographic findings.

The summarization of the differential diagnoses was not part of the program's diagnostic reasoning and was developed for the evaluation to reduce the workload for the reviewers. Indeed, it decreased the number of nodes in the hypotheses to a manageable number, making an evaluation of the detailed reasoning feasible. Because it proved to be a useful tool to convey the important conclusions in the diagnosis, we intend to incorporate it into the program. At present, the summarization can also obscure the nature of relationships present in the hypothesis. The main problem is that not all of the summary links have the same meaning. The labels need to reflect important distinctions such as *causing* versus *possibly contributing to* or *chronic with acute worsening*. Using the word *causes* for all links often misrepresents the kind of relationship that exists. There are also a number of conventions used by people in conveying degrees of uncertainty that need to be incorporated in the summaries, such as saying *pulmonary disease* rather than *COPD* when the findings are nonspecific. Thus, summarization of the differential diagnoses is difficult to do well and more work is needed to bring it to the competence necessary to be an effective utility.

Evaluation Methodology

Evaluation of expert systems is always difficult, but it was more difficult for the HFP because of the complex conclusions that were being evaluated. Conducting the evaluation as a single session with all reviewers was extremely beneficial. It enabled us to

answer their questions as they arose, monitor their progress, make sure they understood the instructions as intended, and encourage them to provide comments with their judgments. It also allowed for some discussions, which slowed the progress. Since three hours is about the maximum length of time people can effectively do such detailed work, a larger evaluation may require multiple sessions or the rest of the reviews to be done in a different context.

The use of evaluation forms proved to be much more practical than using the program interactively because of the logistics involved, the additional delays that the program would have introduced, and the ability of paper to take comments anywhere. The evaluation forms themselves left us with the difficult task of analyzing the judgments to determine the source of criticisms. Fortunately, the reviewers wrote enough comments to make that analysis possible. The problem is essentially one of designing a multiple-choice-question test when an essay is needed. The reviewers are being asked to critique the details, but there are too many details to consider them one by one. The node interconnections imply that a single issue changes the rating of multiple nodes. A possible improvement would be to organize the questions around paths through the summaries rather than single nodes, or to organize the questions around "diagnostic" nodes (the main disease and syndrome nodes).

The automatic generation of the evaluation forms utilized a limited vocabulary and "computer logic" that made some of the statements difficult to understand. Particularly difficult were statements about nodes appearing only in some of the hypotheses. The significantly poorer rating of such possible nodes and the smaller number of correct ratings of alternative hypotheses may have been due to these problems of presentation or it may be that determining the range of a differential is a more difficult problem than determining the best hypothesis.

The final issue that makes an evaluation of diagnostic programs difficult is the lack of a "gold standard." It is tempting to say that the final diagnosis of the patient is the gold standard, but the objective of the differential is to determine all of the hypotheses that are consistent with the patient presentation rather than only the diseases that the patient actually has. Because there is no objective way of obtaining such a list, the best that can be achieved is expert consensus. As a result, there will always be some controversy and misunderstanding in the critiques. A common strategy to control for these factors is to have the experts critique each other as well as the program.

That is not feasible when the diagnoses are detailed because the expert diagnoses are given with much more selective details. An alternative would be to have the reviewers individually critique a case and then collectively agree on a final critique. The problem, of course, is that it would lengthen an already time-consuming and somewhat tedious process.

Even with the difficulties of designing and analyzing this type of evaluation, the evaluation has been very helpful in determining the level of performance of the program. Overall, the program is capable of providing high-quality detailed diagnostic hypotheses for complex cardiovascular cases. With some additional refinement of the knowledge base and processing of the hypotheses, the mistakes encountered by these reviewers should be eliminated and the error rate decreased significantly. Once these changes and the summarizer have been appropriately validated, the appropriate next step is a prospective evaluation to address the usefulness of the program.

References ■

1. Long WJ, Naimi S, Criscitiello MG, Larsen G. Differential diagnosis generation from a causal network with probabilities. In: Proceedings, Computers in Cardiology Conference. Washington, DC, 1988:185-8.
2. Long W. Medical diagnosis using a probabilistic causal network. *Appl Artif Intell*. 1989;3:367-83.
3. Long WJ. Flexible reasoning about patient management using multiple models. *Artif Intell Med*. 1991;3:3-20.
4. Kingsland LC. The evaluation of medical expert systems: experience with the AI/RHEUM knowledge-based consultant system in rheumatology. In: Ninth Symposium on Computer Applications in Medical Care, Washington, DC, 1985:292-5.
5. Miller RA, Pople HE, Myers JD. INTERNIST-I, an experimental computer-based diagnostic consultant for general internal medicine. *N Engl J Med*. 1982;307:468-76.
6. Feldman MJ, Barnett GO. An approach to evaluating the accuracy of DXplain. In: Fourteenth Symposium on Computer Applications in Medical Care, Washington, DC, 1990:38-43.
7. Long WJ, Naimi S, Criscitiello MG. Development of a knowledge base for diagnostic reasoning in cardiology. *Comput Biomed Res*. 1992;25:292-311.
8. Long WJ. The probability of disease. In: Fifteenth Symposium on Computer Applications in Medical Care, Washington, DC, 1991:619-23.
9. Kahn CE. Validation, clinical trial, and evaluation of a radiology expert system. *Meth Inform Med*. 1991;30:268-74.
10. van der Lei J, Musen MA, van der Does E, Manintveld AJ. Critiquing physician decision making using data from automated medical records: assessing the limitations. In: Fourteenth Symposium on Computer Applications in Medical Care, Washington, DC, 1990:559-63.
11. Yu VL, Fagan LM, Wraith SM, et al. Antimicrobial selection by a computer: a blinded evaluation by infectious diseases experts. *JAMA*. 1979;242:1279-82.
12. Hickam DH, Shortliffe EH, Bischoff MB, Scott AC, Jacobs CD. The treatment advice of a computer-based cancer chemotherapy protocol advisor. *Ann Intern Med*. 1985;103:928-36.

APPENDIX

Example Evaluation Form

Note: Annotations that were not part of the actual form are in bold. The following is the input description given to the program, to allow reviewers to consider their own diagnoses:

Patient: PT2006

HISTORY: 65-year-old male with normal weight, orthopnea progressive, dyspnea on-exertion progressive and cough, having known diagnoses of congestive cardiomyopathy, atrial fibrillation, chronic renal insufficiency and diabetes and on furosemide, nitroglycerin, captopril, digitalis, dobutamine for 1 wk and dopamine for 1 wk

VITAL SIGNS: bp: 114/76 hr: 77, R: 18 and T: 97.4

PHYSICAL EXAM: appears in no acute distress, conscious, chest revealed basilar rales, jugular pulse shows jvp: 12 cmH₂O, pulse is irregular, auscultation revealed fixed splitting s₂, lv_s₃, I-II/VI holosystolic systolic murmur at the apex and III/VI holosystolic systolic murmur at the llb, laterally displaced apex impulse, ascites and mild pedal edema

LABORATORY FINDINGS: ekg (1 hr): atrial fibrillation, cxr (1 hr): congestive failure, electrolytes (1 hr): Na:134 k:4.4 bun:69 creat:2.2 Ca:10.0 serum albumin:4.4, blood gases (1 hr): po₂:83 ph:7.49 pco₂:36, cbc (1 hr): hct:39 wbc:9.3, increased cpk-mb (1 day) and elevated liver function tests (1 hr)

HEMODYNAMIC DATA: w:30, ci:2.3, pap:54/24, rap:20 and data from swan-line (now)

New page, so the reviewer is less influenced by the program's conclusions.

This is a graphic summary of the leading assessment generated by the program, followed by a textual description and the differences in other hypotheses. The following pages will ask for a critique of the parts of this and the rest of the hypotheses.

- CONGESTIVE FAILURE caused by ATRIAL FIBRILLATION, CONGESTIVE CARDIOMYOPATHY, DOPAMINE TOXIC accounts for cough, orthopnea, dyspnea on exertion, basilar rales, jvp:12 cmH₂O, ascites, R:18, mild pedal edema, hct:39, Na:134, ci:2.3, w:30, cxr:congestive failure, rap:20, elevated liver function tests, bun/creat:31, creat:2.2, bun:69, HIGH PA PRESS, FUROSEMIDE TOXIC

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- CONGESTIVE CARDIOMYOPATHY accounts for known diagnoses of congestive cardiomyopathy, laterally displaced apex impulse, lv s3, CONGESTIVE FAILURE, DOBUTAMINE, MITRAL REGURGITATION, DOPAMINE TOXIC, CAPTOPRIL, and possibly, DIGITALIS

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

"Possibly digitalis" means one hypothesis made this causal connection and the other did not.

- ATRIAL FIBRILLATION caused by MITRAL REGURGITATION accounts for known diagnoses of atrial fibrillation, pulse is irregular, DIGITALIS

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- MITRAL REGURGITATION caused by CONGESTIVE CARDIOMYOPATHY accounts for holosystolic systolic murmur at the apex, holosystolic systolic murmur at the llsb, ATRIAL FIBRILLATION

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- RENAL INSUFFICIENCY accounts for known diagnoses of chronic renal insufficiency, creat:2.2

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- HIGH PA PRESS caused by CONGESTIVE FAILURE accounts for mean pap:34

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- NITROGLYCERIN treating possibly, CORONARY ARTERY DISEASE

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- FUROSEMIDE TOXIC treating CONGESTIVE FAILURE accounts for mixed alkalosis

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

"Furosemide toxic" just means that the hypothesis is attributing side effects to the drug.

- DOPAMINE TOXIC treating CONGESTIVE CARDIOMYOPATHY accounts for CONGESTIVE FAILURE

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- DIGITALIS treating ATRIAL FIBRILLATION, and possibly, CONGESTIVE CARDIOMYOPATHY

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

Both of the following are possibly unaccounted since each is accounted for in one of the two hypotheses.

- Unaccounted Findings: possibly, fixed splitting s2, increased cpk mb

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

The following are included in some hypotheses:

The first two were only in the first hypothesis, the third node was in the second hypothesis.

- MYOCARDIAL INFARCTION ACUTE caused by CORONARY ARTERY DISEASE accounts for increased cpk mb

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- CORONARY ARTERY DISEASE accounts for MYOCARDIAL INFARCTION ACUTE, NITROGLYCERIN

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- ATRIAL SEPTAL DEFECT accounts for fixed splitting s2

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

Relationships left out: _____

Following is the differential diagnosis of the program. Please critique the hypothesis and summary of alternates in the boxes provided.

These are the same as before, only now the reviewer is asked to critique the best hypothesis and alternates.

- Best Hypothesis

- CONGESTIVE CARDIOMYOPATHY accounts for lv s3, laterally displaced apex impulse, known diagnoses of congestive cardiomyopathy, treated with CAPTOPRIL, DOPAMINE TOXIC, DOBUTAMINE, causing:
 1. MITRAL REGURGITATION accounts for holosystolic systolic murmur at the llsb, holosystolic systolic murmur at the apex, causing:
 - ATRIAL FIBRILLATION accounts for pulse is irregular, known diagnoses of atrial fibrillation, treated with DIGITALIS,
 2. CONGESTIVE FAILURE accounts for bun:69, creat:2.2, bun/creat:31, elevated liver function tests, rap:20, cxr:congestive failure, w:30, ci:2.3, Na:134, hct:39, mild pedal edema, R:18, ascites, jvp:12 cmH2O, basilar rales, dyspnea on exertion, orthopnea, cough, treated with FUROSEMIDE TOXIC, causing:
 - HIGH PA PRESS accounts for mean pap:34,
- CORONARY ARTERY DISEASE treated with NITROGLYCERIN, causing:
 - MYOCARDIAL INFARCTION ACUTE accounts for increased cpk mb,
- RENAL INSUFFICIENCY accounts for creat:2.2, known diagnoses of chronic renal insufficiency,
- left unaccounted: fixed splitting s2,

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

- Additional Hypotheses may add ATRIAL SEPTAL DEFECT exclude MYOCARDIAL INFARCTION ACUTE, CORONARY ARTERY DISEASE

correct	possible	partly correct	wrong	seriously wrong	comment:
---------	----------	----------------	-------	-----------------	----------

The following was used by the reviewers to indicate alternatives that they felt were left out.

Missing hypotheses: _____

Comments: _____

In addition a final page was included with the detailed graphic representation of the best hypothesis, in case the summary representation was insufficient.