Heuristic Abstraction in the Decision-Theoretic Pathfinder System*

Eric J. Horvitz
Medical Computer Science Group
Knowledge Systems Laboratory
Departments of Computer Science and Medicine
Stanford, California 94305

David E. Heckerman
Medical Computer Science Group
Knowledge Systems Laboratory
Departments of Computer Science and Medicine
Stanford, California 94305

Keung-Chi Ng
Computer Science Department
University of Southern California
Los Angeles, California 90089

Bharat N. Nathwani
Department of Pathology
School of Medicine
University of Southern California
Los Angeles, California 90303

*Core research on the Pathfinder project has been supported by the National Library of Medicine under Grant RO1LM04529. Components of the research have been supported by the by the National Aeronautics and Space Administration under Grant NCC-220-51, by the National Science Foundation under Grant IRI-8703710, and by the Josiah Macy, Jr. Foundation. Computational support has been provided by the SUMEX-AIM Resource under NIH Grant RR-00785.

Abstract

A criticism of diagnostic systems, which are based on the formal foundations of probability and utility, is that their reasoning strategies and recommendations are inflexible and unnatural. We have developed a facility that increases the flexibility of normative reasoning systems by providing multiple human-oriented perspectives on diagnostic problem solving. The method endows a system with the ability to reason about arbitrary classes of diagnostic entities and to control the level of abstraction at which inference occurs. The techniques have been integrated into Pathfinder, an expert system that performs hematopathology diagnosis. We explain the background and approach that we have taken, and describe how we use the techniques in Pathfinder to modulate information- and decision-theoretic reasoning with strategic scripts that are familiar to physicians.

Introduction

Surveys of the preferences of clinicians have identified the importance of the understandability of the reasoning of an expert system as an important factor in its acceptance [27,3]. The transparency of inference has been considered a definitive component of expert systems, distinguishing them from numerical programs and other kinds of reasoning systems in artificial intelligence [2]. The important role of reasoning transparency in expert systems has made explanation an important area of artificial-intelligence research [24].

A valid criticism of early diagnostic systems based on the formal foundations of probability and decision is that clinicians found their reasoning strategies and recommendations unnatural and difficult to understand [26,5,23]. Problems with the expressiveness and flexibility of these normative systems have been a motivation for the investigation of quasi-probabilistic and ad hoc approaches for reasoning under uncertainty [9].

The Pathfinder team has pursued the solution of theoretical and pragmatic problems with the construction and effective use of a medical expert system for hematopathology diagnosis based on the principles of decision-theory [12]. In this paper, we describe work on strategic reasoning in Pathfinder that has been motivated by problems with the inflexibility and opacity of straightforward implementations of normative problem solving.

Several investigators have addressed problems with understanding normative reasoning have developed facilities for qualitatively explaining the results of probabilistic and decision-theoretic inference [19,18,15]. We have worked on generating more natural decision-theoretic inferences and explanations by allowing a clinician to control the level of abstraction at which inference occurs. For example, rather than directly reasoning about the beliefs associated with each disease, in response to evidence observed in a tissue section under a microscope, a pathologist may prefer to reason about classes of disease, such as inflammatory, infectious, and malignant. At this higher level of abstraction, the uncertain reasoning problem is simplified, and thus is easier to understand and explain. The availability of a set of such heuristic abstraction strategies allows a clinician to probe a diagnostic problem from a variety of familiar perspectives.

The Hematopathology Problem Area

The Pathfinder project was initiated to solve problems that general pathologists have making hematopathology diagnoses. In particular, we have worked to build a computer-based assistant to guide pathologists in the interpretation of several hundred histologic features that appear in sections of lymph-node tissue. The microscopic interpretation of lymph-node biopsies has been considered to be one of the most difficult and error-prone tasks of surgical pathology. The complexity of lymph-node pathology has led to major problems in the diagnosis of lymph-node diseases. Many of the benign diseases of the lymph node closely resemble the malignant diseases. The accurate diagnosis of diseases that present as complex visual patterns in lymph-node tissue is crucial for the determination of prognosis and therapy. Malignant lymphomas typically have characteristic responses to therapy and different survival rates. A computer-based decision-support system could be useful in bringing expert knowledge and experience to the general pathologists, and help reduce the difference in quality between the diagnoses made at community hospitals and those made by the handful of experts that specialize in lymph-node diagnoses.

Computational Architecture of Pathfinder

We constructed several different implementations of the Pathfinder expert system. The earliest Pathfinder expert system was developed on the Sumex-AIM DEC-2060 in the MRS logic-programming language. Preliminary experimentation with this system highlighted the difficulty of representing and manipulating the uncertain relations between features and diseases. We implemented several quasi-probabilistic methods similar to the approach used in the Internist-1[22] and QMR[21] systems before moving to probability theory. That move was motivated by the precise definition of probability and by significant increases in the diagnostic accuracy of the system. Later versions of Pathfinder were implemented on the HP-9836 workstation. More recently, the system was reimplemented in object Pascal within the MPW environment on the Macintosh II. The current system reasons about approximately 60 malignant and benign diseases of lymph-nodes, constructing differential diagnoses through the consideration of evidence about the status of up to 100 morphologic features presenting in lymph-node tissue.

The computational architecture of the Pathfinder system is based on the *hypothetico-deductive* approach to diagnosis (also called the *method of sequential diagnosis* in the medical informatics literature[10,9]). A flow-chart representation of this method is shown in Figure 1. A set of salient disease manifestations are initially presented to the program by the pathologist. A differential diagnosis is then formulated based on these manifestations, and questions are selected that can help to decrease the number of diseases under consideration. After the user answers one or more of these questions, a new set of hypotheses is formulated; the process is repeated until a diagnosis is reached. In the following sections, we will discuss the details of several of these steps.

In Pathfinder, *features* are each structured into a set of 2 to 10 mutually exclusive and exhaustive list of *values*. For example, the feature *pseudofollicularity* can take on any one of the values *absent*, *slight*, *moderate*, or *prominent*. A pathologist evaluates features by

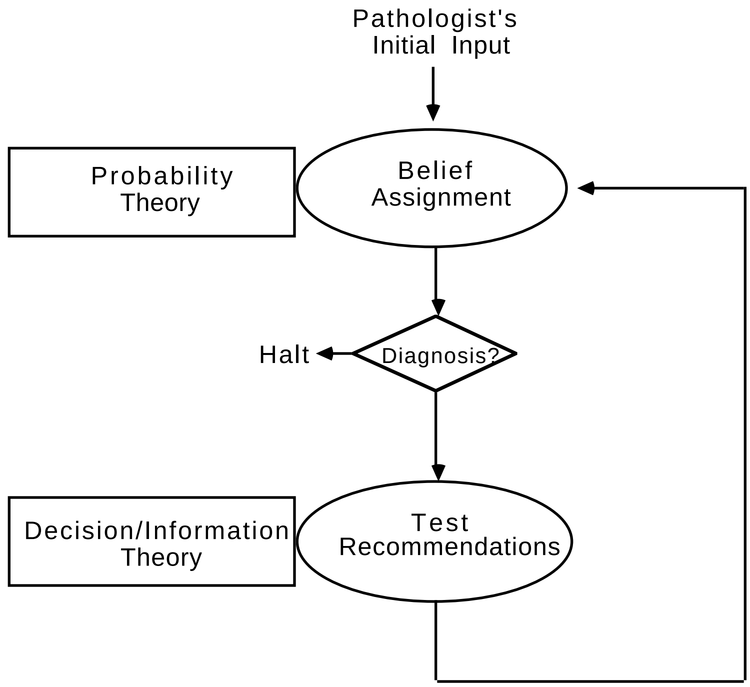

Figure 1: The Pathfinder hypothetico-deductive architecture. Initial evidence presented to the system is used to form a list of hypotheses. Next, a value-of-information analysis identifies the next best tests.

selecting a value that reflects the severity of a feature, forming a *feature–value* pair.

Value of Information in Pathfinder

In practice, the Pathfinder system allows a user to enter information about microscopic observations from a biopsy section. A probabilistic reasoner apportions a level of belief to alternative diseases in the system; after this phase, the system applies a value-of-information analysis to make new test recommendations.

As each feature consists of a set of mutually exclusive and exhaustive values, we can denote the possible evidence associated with a particular feature, F, by $F_1, F_2, \ldots, F_n$, where n is the number of mutually exclusive values associated with the feature.

Original versions of Pathfinder used information-theoretic measures for question selection. In pathology—where pathologists tend to continue their diagnostic analysis until they have ruled out all but a single disease, or have given up and requested another section—the informational approach is equivalent to the assumption of equal cost for all tests and of equal cost for all misdiagnoses. Although this assumption is valid from the point of view of pathologists, it would clearly be suboptimal for patient care in situations of incomplete diagnosis and high test costs. Thus, later versions of Pathfinder make available full expected-utility optimization of value-of-information. We have assessed a full disease utility model in Pathfinder, making use of the work of Howard on life and death decision making [16,12].

We will now present the information-theoretic strategy. This approach selects features that give the *lowest* expected entropy, $H(DD, F)$ associated with a feature F in the context of the current differential diagnosis, DD,

$$H(DD, F) = \sum_i p(F_i)\ H(DD, F_i)$$

The quantity $H(DD, F_i)$ denotes the entropy of the differential when F_i is observed, and is given by

$$H(DD, F_i) = -\sum_j p(D_j|F_i)\ \log p(D_j|F_i)$$

The quantity $p(F_i)$ is calculated using the probability expansion rule

$$p(F_i) = \sum_j p(F_i|D_j)\ p(D_j)$$

where $p(D_j)$ refers to the probability of disease on the differential diagnosis before F_i is observed.

Abstraction in Hematopathology

Early work by our expert-systems research group uncovered a set of inference-related issues that we attributed to human cognitive-resource constraints and preferences [11,13]. While building and testing Pathfinder, we found that straightforward applications of decision-theoretic inference could produce behaviors that were seen as counterintuitive and confusing. We found that the information- and utility-maximizing strategies, used in determining the next set of features for a pathologist to evaluate, performed analyses at a more detailed level than those with which some system users were familiar with. That is, our early, inflexible versions of Pathfinder represented and reasoned only about the finest diagnostic distinctions in hematopathology. We found that users preferred to work at higher levels of abstraction than those used in our inflexible decision-theoretic approach. We also noticed that users had different preferences about transitions from one type of abstraction to another.

Discussions with users unearthed alternative problem-solving hierarchies that often were used to segment a single complex diagnostic-reasoning task (from the perspective of the decision-theoretic system) into a set of tasks at increasingly detailed levels of abstraction. These human-oriented abstraction strategies allow a pathologist to reason about discriminating among groups of similar diseases, rather than to consider each disease as a separate entity.

As an example, we found that the expert pathologist on the Pathfinder team often imposes a simple two-group discrimination structure on the problem-solving task. As opposed to a strategy of discriminating among all the diseases on the differential diagnosis, the pathologist's discrimination task at any point in reasoning about a case is constrained to a small number of groups of diseases. As categories of diseases are ruled out, the particular pairs of groups considered become increasingly specific. For example, if there are benign and malignant diseases on a differential diagnosis, the pathology expert often deems most appropriate those questions that best discriminate between the benign and malignant groups, rather than questions that might best discriminate among all possible diseases. If all benign diseases have been ruled out, leaving only primary malignancies and metastatic diseases on

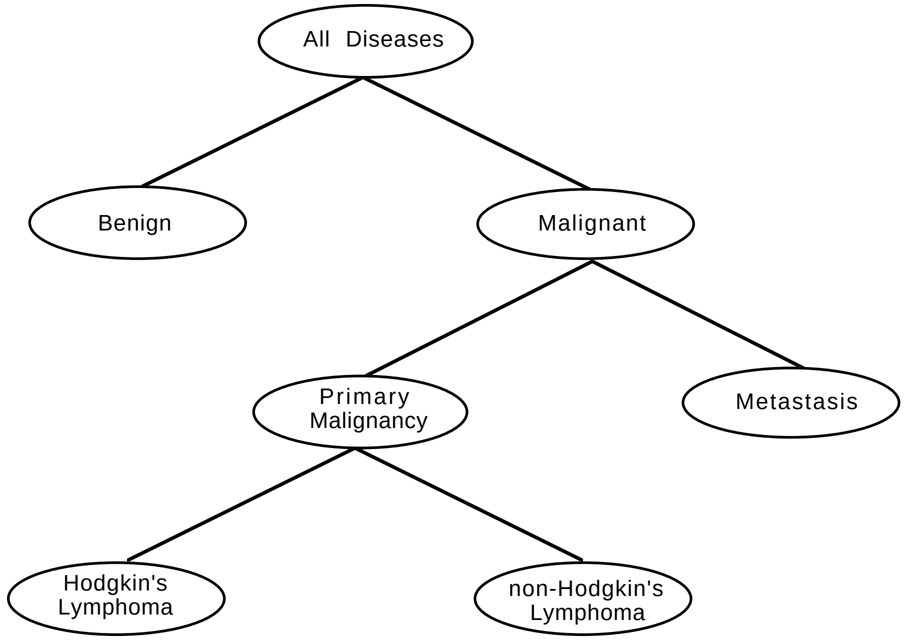

Figure 2: A heuristic problem-solving hierarchy showing how a pathologist may categorize diseases into a sequence of abstraction classes to manage the complexity of diagnostic inference.

the differential diagnosis, the pathologist will attempt to discriminate between the primary malignancy and the metastatic categories.

The expert's diagnostic strategy was often described by the traversal of a hierarchy of fundamental categories of disease. The problem-solving hierarchy shown in Figure 2 is a binary tree of disease groups. The hierarchy can be used to group the differential diagnosis at various levels of refinement.

We set out to develop techniques for allowing a user to request that a decision-theoretic–based expert system reason more "naturally" by *constraining* value-of-information calculations to a set of alternative preferred abstractions, as well as to particular preferred transitions between the subtasks. In general, we can show that each of these subproblems is less optimal than is the single best step computed by the automated complex analysis. Nevertheless, users may desire to receive assistance, and to learn about diagnosis, from the perspective of preferred abstractions.

Previous Studies of Complexity and Abstraction

Research by cognitive psychologists have addressed issues that have relevance to the the use of abstraction in computer-based decision systems for clinicians. Cognitive-psychology experiments have demonstrated severe limitations in the ability of humans to consider more than a handful of concepts in the short term [17,8]. One cognitive study revealed that humans cannot retain and reason about more than two concepts in an environment with distractions [28]. Abstracting a set of atomic entities into larger classes enables a person to reduce the number of entities under consideration. Psychologists have speculated that hierarchies of abstraction are used often by people to facilitate easy indexing through simple relationships among classes at different levels of abstraction [20,25]. Indeed, cognitive psychologists have found that humans make use of abstraction hierarchies in a variety of domains. Of particular relevance to our work, psychologists studying medical decision making have found that physicians in specialty areas of medicine frequently make use of abstraction strategies for managing the complexity of clinical problem solving [6,7].

Previous work in artificial intelligence has addressed the usefulness of heuristic classification strategies in diagnosis in rule-based systems. Most notably, the work by Clancey on the Neomycin system focused on the application of strategic knowledge about useful abstractions to control reasoning [4]. Ben-Bassat described the need for tools that allow a user to specify abstract diagnostic classes in probabilistic systems[1].

Heuristic Abstraction in Pathfinder

The discovery of the use of an abstraction hierarchy by our domain expert suggested the development of a new question-selection strategy that could discriminate among classes of diseases, instead of among individual diseases. We hoped that design and application of such a strategy would make explanation clear because the user would have to consider the relevance of a recommendation to only two or three groups.

Our attempt to constrain the discriminatory focus of the evidence-gathering strategy led to a new reasoning strategy [11]. This *group-discrimination* strategy selects questions based on their ability to discriminate between disease classes contained at a single level of the abstraction hierarchy. As diseases are ruled out, the system moves progressively to more specific levels of the hierarchy.

For a given differential diagnosis, the group-discrimination strategy identifies the most specific grouping possible, then selects questions that best discriminate among groups of diseases.

More formally, suppose the differential is split into two groups, G_a and G_b, of n_a and n_b diseases, respectively:

$$G_a = (D_{a1}, D_{a2}, \ldots, D_{an_a})$$

$$G_b = (D_{b1}, D_{b2}, \ldots, D_{bn_b})$$

As we assume that only one lymph-node disease is present in Pathfinder, we can consider the diseases to be mutually exclusive events. For each group, we are interested in the probability that the true diagnosis will be in that group. To calculate this probability, we sum the probabilities of all the diseases within each group. That is, the probability that a group contains the true diagnosis, given F_i is

$$p(G_j|F_i) = \sum_k p(D_{jk}|F_i) \qquad j = a, b$$

Therefore, the expected entropy of the grouped differential, denoted $H_G(DD, F)$, is given by

$$H_G(DD, F) = \sum_i \; p(F_i) \; H_G(DD, F_i)$$

where $H_G(DD, F_i)$ is the entropy of the grouped differential given feature–value F_i,

$$H_G(DD, F_i) = -\sum_j p(G_j|F_i)\ \log p(G_j|F_i)$$

The expected entropy, $H_G(DD, F)$, is inversely related to the additional information contained in F_i about the grouped differential diagnosis. The group-discrimination strategy selects those features that minimize $H_G(DD, F)$.

The grouping strategy ignores information concerning the probabilities of diseases within each group. Only the probability that the true diagnosis lies within a group is considered in the calculations. Because the grouping strategy has a simpler discriminatory focus and more closely follows the decision-making protocol of the expert lymph-node pathologist than did the initial information-optimizing strategy, it is relatively easy to explain. Instead of having to present complex summaries explaining how each piece of evidence might affect belief in the presence of a number of diseases, an explanation of questions generated by group-discrimination strategy must simply demonstrate how possible responses affect the two groups under consideration.

A predictable problem with the use of the group-discrimination strategy is that the differential-diagnosis–refinement process will not always proceed as quickly as it does with the application of the more optimal strategies. That is, the grouping value-of-information strategies are not as efficient as are the more powerful nongrouped strategies; on average, we expect a larger number of evidence-gathering requests will be made by the group-discriminate strategies to achieve a similarly refined differential diagnosis (in a utility or informational sense). We must expect a larger number of feature identifications, because detailed information about the relative plausibility of individual diseases within each group is discarded in the grouping process.

Generating Multiple Perspectives with Abstraction

Further research revealed additional richness in the heuristic complexity-management schemes used by hematopathologists . There are a number of alternative *perspectives* on any given diagnostic problem; that is, there are several ways to manage the complexity of diagnostic reasoning. A base differential diagnosis can be reformulated heuristically into disease classes at varying levels of abstraction. The alternative formulations often reflect perspectives on the pathological entities considered to have high discriminatory value.

Figure 3 shows a classification strategy based on the *origin* of the dominating population of proliferating cells. This strategy is used at times by the chief expert on the project. Figure 4 shows a decomposition strategy employed by a hematopathology resident who singled out by our domain expert as having a special gift for the diagnosis of lymph-node pathology. Notice that this grouping strategy stresses a high-level gestalt approach to *patterns* of morphological pathology. Beyond the hierarchies, there is a large number of simple differential-dependent grouping strategies that may be imposed at various times. These schemes are currently stored in a library of complexity-management strategies.

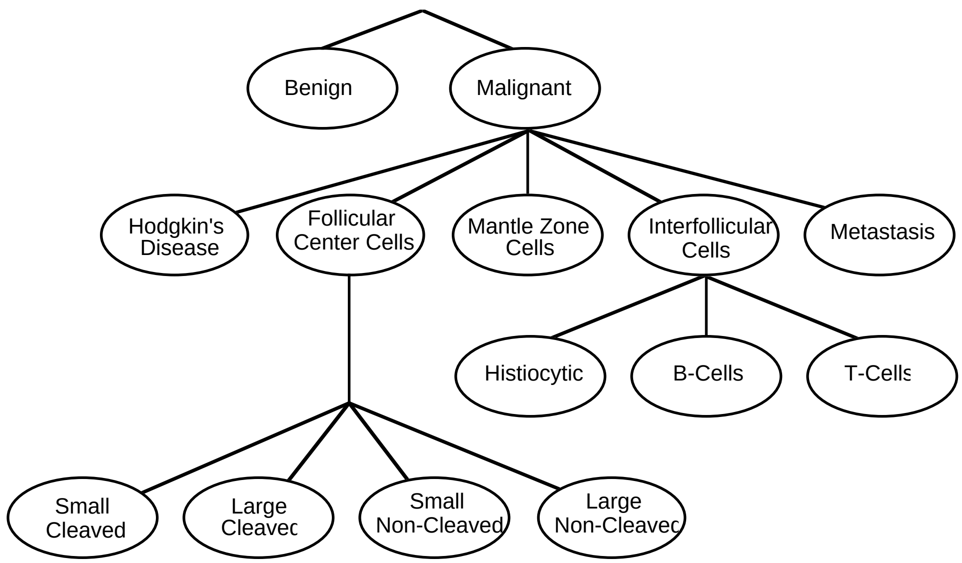

Figure 3: A strategic hierarchy representing the formulation of the diagnostic problem from the perspective of the origin of the predominant proliferating cell line in a lymph-node section.

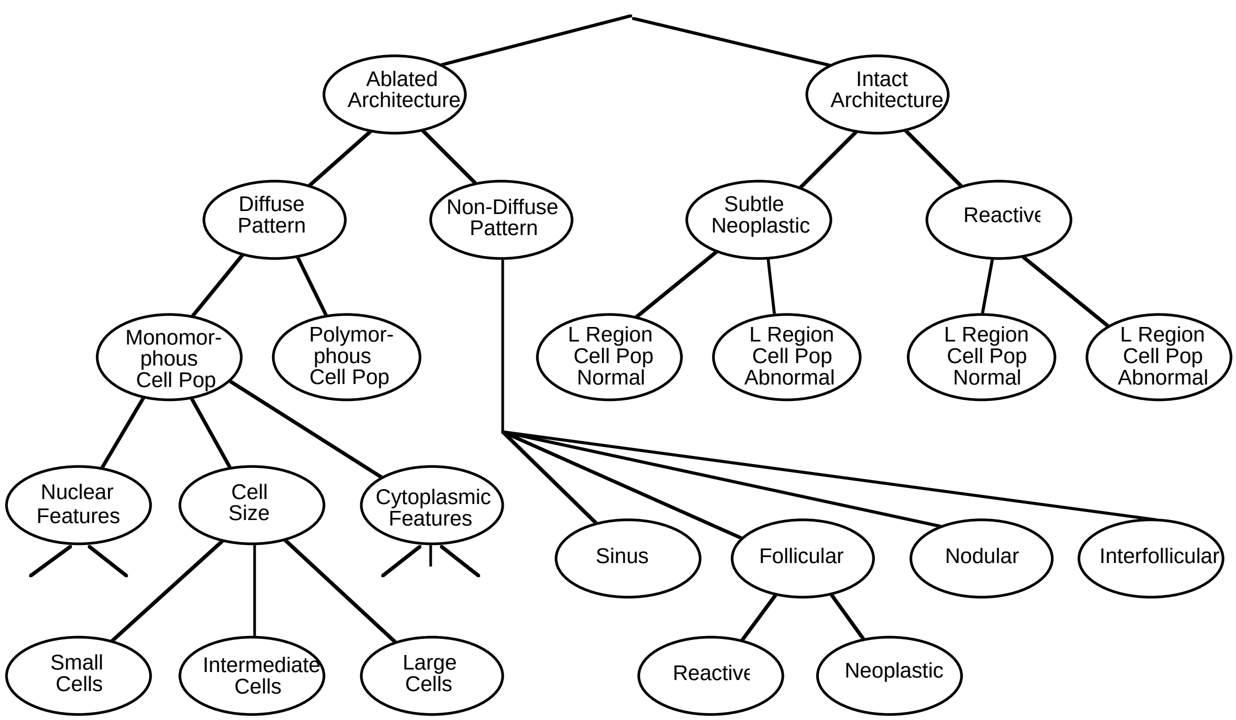

Figure 4: A strategic hierarchy representing the formulation of the diagnostic problem from the perspective of the morphological pattern seen in a lymph-node section.

Heuristic-Abstraction Facility

We have developed a heuristic abstraction facility, based on the grouping strategy described in Section 4, and have integrated it with the Pathfinder system. The facility allows system developers or users to specify a library of intuitive classes of diseases and to nest these abstractions within arbitrary strategic scripts. The classifications and scripts may be invoked at any time during a diagnostic session to guide Pathfinder's recommendations and display. All strategies are available on a pull-down *problem formulation* menu. Invoking a strategic script brings up a window that displays the strategy. The window also displays the differential diagnosis and highlights the current level of abstraction. Probabilities of categories and of individual diseases are listed. Multiple windows–each representing a different perspective on the same problem–may be displayed simultaneously. By clicking on one of the windows, a user activates the perspective. She then can ask the system to generate a recommendation about the best way to discriminate among the leading disease classes represented in the window. When new information becomes available, all windows are updated.

Figure 5 contains a Pathfinder screen showing a current problem from an *etiological* and *pattern* perspective. In the foreground, another perspective capturing the strategy of *leading verses others* is displayed. This strategy will find those tests and features that will rule out contenders for the leading disease. The system allows a user to easily compare questions generated by information- or complete utility-theory approach within the context of any of the human-oriented strategies.

Direction of Research

We are interested in extending the abstraction facility with techniques for using strategic knowledge to automate the selection of perspective for guiding decision-theoretic inference [14]. This approach requires the assessment of multiple attributes of preference about abstraction. The informational costs associated with the use of an abstraction strategy relative to a complete analysis can be calculated dynamically and traded off with benefits based in the simplicity and understandability of reasoning. We are also exploring the costs attributed by clinicians to constraints on the transition from one perspective or level of abstraction to another, as well as from one information class to another (e.g. being asked to move between the evaluation of microscopic features at low- and high- objective at to another). We have done preliminary exploration of the preferences that might be represented in a strategic utility model or in a set of rules about abstraction. In the context of pathology, we desire the strategy selection mechanism to perform as naturally as is possible within some maximum allowed distance from the optimal nongrouped strategy. Another area of research is to enrich the notion of perspective not only to capture preferences about problem abstraction, but also to take into consideration different perspectives of diagnostic utility. For example, a physician may wish to compare the differences between diagnostic recommendations in light of monetary constraints on the tests that are currently available in an institution. We could achieve such flexibility by allowing a clinician to select among different utility models. Finally, we plan on developing a tool that will allow a user to inspect useful abstractions of a current differential diagnosis for any feature.

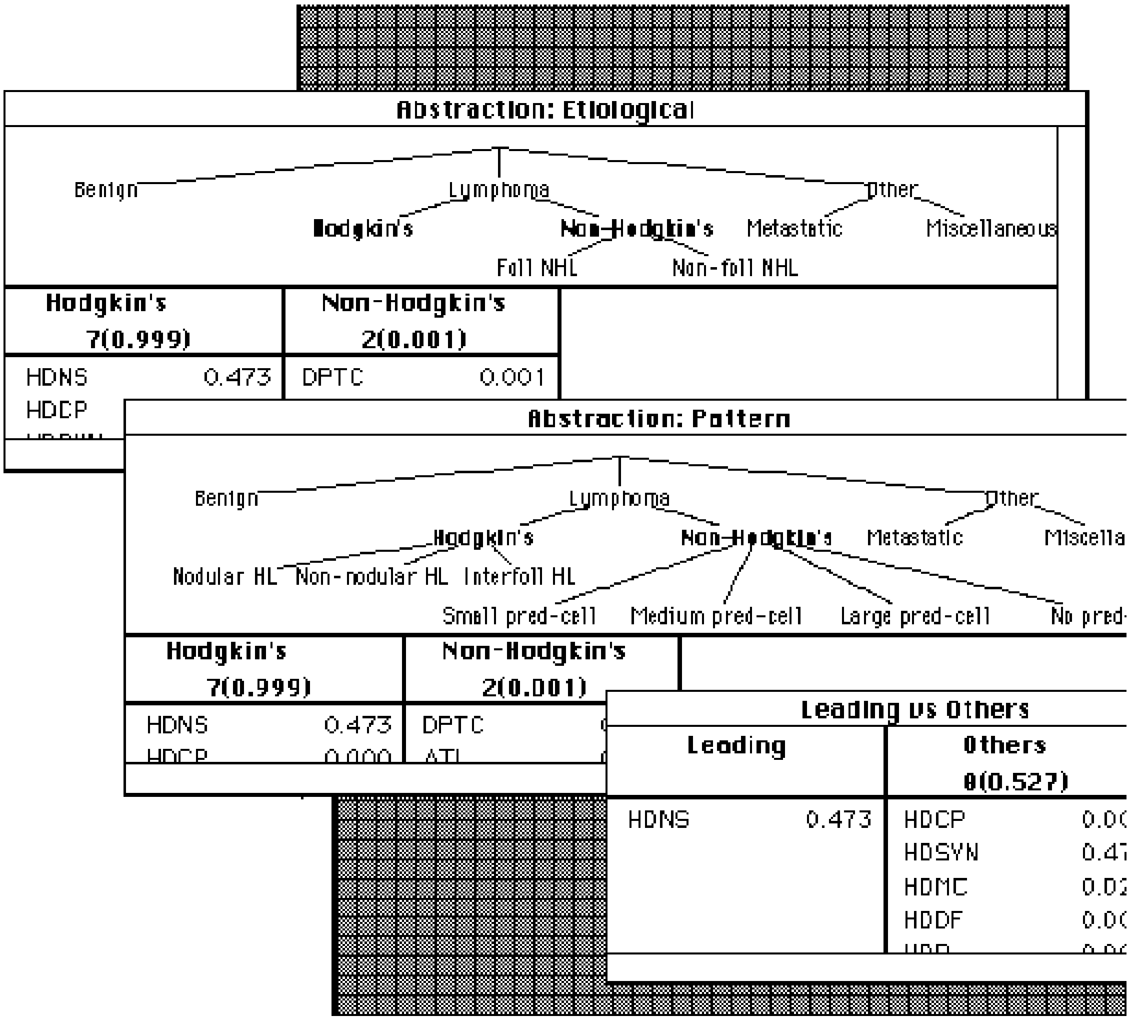

Figure 5: A screen from Pathfinder demonstrating alternative abstraction strategies for managing the complexity of diagnostic inference. A pathologist invokes the system to generate a recommendation tailored to a particular perspective by pointing the mouse cursor at one of the windows. The current active process, in the foreground, displays the problem from a *leading disease verses others* perspective.

Conclusion

The complexity of decision-theoretic inference and the inflexibility of many decision-theoretic diagnostic systems have bolstered the stereotype of such systems as necessarily rigid and unnatural. We have developed techniques that allow us to represent and use heuristic, human-oriented abstractions within a decision-theoretic system. The new capability increases the naturalness of normative inference and explanation. We have found the graceful integration of such flexibility to be useful in adapting decision-theoretic inference to human users. We are continuing to enrich the cognitive style of normative reasoning systems so that they are more compatible with system users. We forsee that such human–computer communication research will help us to create decision-theoretic systems that show greater flexibility and responsiveness to the informational needs of clinicians.

Acknowledgments

We thank Gregory Cooper, Douglas Ellison, Lawrence Fagan, Thomas Lincoln, Randolph Miller, Edward Shortliffe, Peter Szolovits, and Clive Taylor for advice on Pathfinder research. Lyn Dupre and Lawrence Fagan provided useful editorial comments on this manuscript.

References

[1] M. Ben-Bassat and D. Teeni. Human-oriented information acquistion in sequential pattern classification: Part 1 - single membership classification. *IEEE Transactions on Systems, Man, and Cybernetics*, 14:131–138, 1984.

[2] B. G. Buchanan. Research on expert systems. In *Machine Intelligence*, volume 10, pages 269–299. Ellis Howard Ltd., Chichester, England, 1982.

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

[4] W.J. Clancey. Heuristic classification. *Artificial Intelligence*, 27(3):289–350, 1985.

[5] R. Davis. Consultation, knowledge acquisition, and instruction. In P. Szolovits, editor, *Artificial Intelligence in Medicine*, pages 57–78. Westview Press, Boulder, CO, 1982.

[6] A.S. Elstein, M.J. Loupe, and J.G. Erdman. An experimental study of medical diagnostic thinking. *Journal of Structural Learning*, 2:45–53, 1971.

[7] A.S. Elstein, L.S. Shulman, and S.A. Sprafka. *Medical Problem Solving: An Analysis of Clinical Reasoning.* Harvard University Press, Cambridge, MA, 1978.

[8] Miller G.A. The magical number seven, plus or minus two. *Psychological Review*, 63:81–97, 1956.

[9] G.A. Gorry. Computer-assisted clinical decision making. *Methods of Information in Medicine*, 12:45–51, 1973.

[10] G.A. Gorry and G.O. Barnett. Experience with a model of sequential diagnosis. *Computers and Biomedical Research*, 1:490–507, 1968.

[11] D.E. Heckerman, E.J. Horvitz, and B.N. Nathwani. Pathfinder research directions. Technical report, Stanford University, October 1985. Knowledge Systems Lab Technical Report KSL-89-64.

[12] D.E. Heckerman, E.J. Horvitz, and B.N. Nathwani. Update on the Pathfinder Project. In *Proceedings of the 13th Symposium on Computer Applications in Medical Care*, October 1989.

[13] E. J. Horvitz, D. E. Heckerman, B. N. Nathwani, and L. M. Fagan. The use of a heuristic problem-solving hierarchy to facilitate the explanation of hypothesis-directed reasoning. In *Proceedings of MEDINFO-86*, pages 27–31, October 1986.

[14] E.J. Horvitz. A multiattribute utility approach to inference understandability and explanation. Technical Report KSL-28-87, Stanford University, Knowledge Systems Laboratory, Stanford, CA, March 1987.

[15] E.J. Horvitz, J.S. Breese, and M. Henrion. Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, 2:247–302, 1988.

[16] R.A. Howard. On making life and death decisions. In R.A. Howard and J.E. Matheson, editors, *Readings on the Principles and Applications of Decision Analysis*, volume II, chapter 26, pages 483–506. Strategic Decisions Group, Menlo Park, CA, 1980.

[17] G.A. Austin J.S. Bruner, J.J. Goodnow. *A Study of Thinking*. Wiley and Sons, 1956.

[18] D. Klein. Explaining and refining decision-theoretic choices. Technical Report MS-CIS-87-57 (LINC 74), Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, 1987.

[19] C. P. Langlotz, E. H. Shortliffe, and L. M. Fagan. A methodology for generating computer-based explanations of decision-theoretic advice. *Medical Decision Making*, 8(4):290–303, 1988.

[20] M.D. Mesarovic, D. Macko, and Y. Takahara. *Theory of Hierarchical, Multilevel Systems*. Academic Press, New York, 1970.

[21] R.A. Miller, M.A. McNeil, S.M. Challinor, F.E. Masarie, and J.D. Myers. The Internist-1/Quick Medical Reference project–status report. *The Western Journal of Medicine*, 145:816–822, 1986.

[22] R.A. Miller, E.P. Pople, and J.D. Myers. Internist-1: An experimental computer-based diagnostic consultant for general internal medicine. *New England Journal of Medicine*, 307:476–486, 1982.

[23] P.E. Politser. Explanations of statistical concepts: Can they penetrate the haze of Bayes? *Methods of Information in Medicine*, 23:99–108, 1984.

[24] E.H. Shortliffe. The computer and medical decision making: Good advice is not enough. *IEEE Engineering in Medicine and Biology Magazine*, 1(2):16–18, 1982. Adapted as Chapter 32 in B.G. Buchanan and E.H. Shortliffe, editors, Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project, Addison-Wesley, Reading, MA,1984.

[25] H.A. Simon. The organization of complex systems. In H.H. Pattee, editor, *Hierarchy Theory*, pages 3–27. G. Braziller, 1973.

[26] P. Szolovits. Artificial intelligence in medicine. In P. Szolovits, editor, *Artificial Intelligence in Medicine*, pages 1–19. Westview Press, Boulder, CO, 1982.

[27] R.L. Teach and E.H. Shortliffe. An analysis of physician attitudes regarding computer-based clinical consultation systems. *Computers and Biomedical Research*, 14:542–558, 1981.

[28] N.C. Waugh and D.A. Norman. Primary memory. *Psychological Review*, 72:89–104, 1965.