

Innovations at the Human—Computer Interface: A Medical-Informatics Perspective*

Eric J. Horvitz

Knowledge Systems Laboratory
Departments of Computer Science and Medicine
Medical School Office Bldg X-215
Stanford, California 94301-5479

Introduction

Research on computer-based decision support in medical informatics has been concerned largely with the development of representation and inference methodologies, and the delivery of reasoning techniques on traditional computing platforms. Nevertheless, there has been growing attention to the development of more powerful and elegant human-computer interfaces, with new interaction modalities, including the use of perspective-following simulation and telepresence. I shall describe several promising directions and opportunities surrounding the human-computer interface in computer-based medical systems. I shall first discuss the form of decision support that might someday be accessed regularly during the practice of medicine. Then, I will describe the utility-directed display of information, intelligent user-modeling, and information sharing. Finally, I will present several new platforms and functionalities that I see playing a significant role in the near future. These include a set of techniques allowing hands-free interaction, methods for overlaying or keying information to anatomic structures during surgery, and micropresence--a technique that allows physicians to effectively enter and perform microsurgery in small areas of a patient's anatomy.

Automated Inference for Medical Decision Support

As background on automated reasoning in medicine, there have been major strides in the development of techniques for automated reasoning about likely disorders and pathophysiologic states, given a set of salient symptoms or test results. Several reasoning systems operate based on the method of hypothetico-deductive reasoning, schematized in Figure 1. With this approach, salient findings are input to a reasoning system. A method for assigning belief to alternative competing disorders is applied, which leads to a rank-ordering of diseases, listed by probability. Next, a method for choosing the next best tests to perform or next observations to make is employed. This method considers the tradeoffs in the costs and benefits of alternative tests. When no

*Proceedings of *Medicine Meets Virtual Reality: Applications for 3-D Multimedia Technology in the Health Sciences*, San Diego, California, June, 1992.

test or observation is available that has informational value greater than its cost, a therapeutic action is recommended. Recent systems have focused on the use of automated probabilistic and decision-theoretic technologies to generate and refine a differential diagnosis, and to choose the next best test to select. Systems have been developed for reasoning about emergency-room medicine, internal medicine, intensive-care medicine, neurology, pathology, ophthalmology, and pulmonary disorders. For an overview of recent developments in this area, see Henrion, et al. 1991 and Horvitz, et al. 1988.

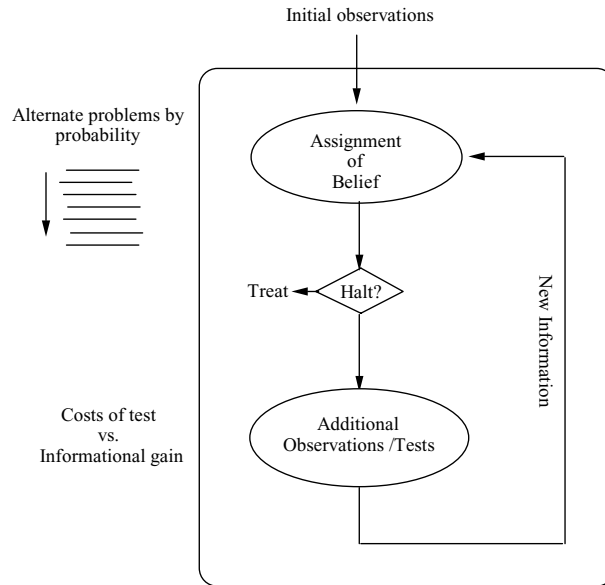


Figure 1. Hypothetico-deductive reasoning in medicine. Cognitive psychologists have found this cycle of inference and testing to be commonplace in clinical reasoning.

Many of the recent systems have been based on related representations called *belief networks* and *influence diagrams*. To construct a belief network, an expert clinician defines nodes that represent important distinctions about the world. These distinctions include hypotheses, important states of interest (for example, diseases) that may not be confirmed directly, and observations that are useful for discriminating among the alternate hypotheses. The expert also provides information about the probabilistic dependencies among nodes.

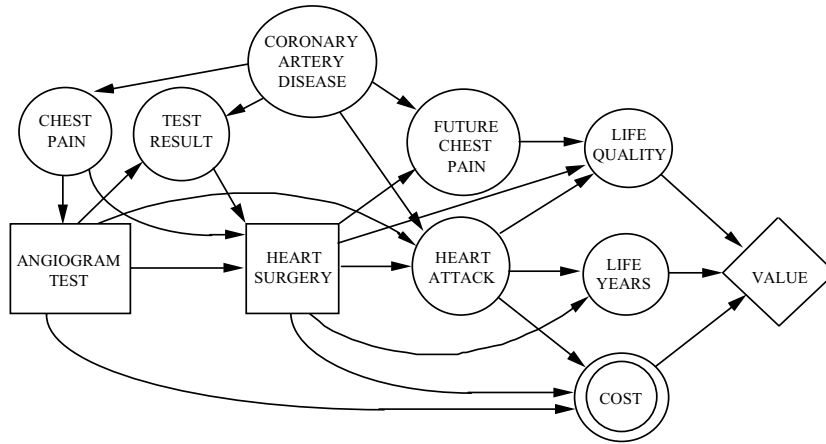


Figure 2. An influence diagram represents a decision problem under uncertainty, with chance nodes (circles), decision nodes (squares), and a value node (diamond). Here, we represent the problem of whether or not a patient should receive an angiogram, given chest pain (from Horvitz, et al. 1988).

An influence diagram is a generalization of a belief network that includes information about decisions and utilities. Figure 3 displays an influence diagram that represents the problem of deciding whether to perform an angiogram test on a patient with chest pain. In determining the expected value of taking, versus not taking, a costly and potentially dangerous test, we consider how different outcomes of the test will affect a decision to have heart surgery. We also consider factors such as how a test, and future surgery, will effect future chest pain, and the possibility of future myocardial infarction. All arcs point into a set of nodes that represent fundamental attributes of value to the patient. The value node represents a function that combines the value of the basic attributes into a scalar utility.

SELECT FEATURE CATEGORY	FEATURES OBSERVED	DIFFERENTIAL DX Grouped
Follicular features	F number: >90%	0.0+ Benign (2)
F number	F density: Back-to-back	0.0+ Florid follic hyperp
F density		0.0+ AIDS
F outside capsule		
F polarity		
F definition		
F size		
F shape		
F mantle zone		
F M2 encircle		
F M2 concentr		
F M2 islands		
F M2 nuclear		
F starry-sky		
F hemorrhages		
F lymphocyte		
F centers atr		
F radially pe		
F protein ppt		
▲ scro		

CRUCIAL FEATURES TO DISTINGUISH GROUPS

F status

Mitotic figures, non-fol1

Monocytoid cell clusters

F polarity

OTHER STRATEGIES

Distinguish all diseases

Pursue leading disease

Cancel

SELECT DISEASES | CRUCIAL FEATURES | INFO | CAROUSEL | REPORT | PRINT | EXIT

Figure 3. A mouse-driven interface to information from inference on a large belief-network. This system, named Intellipath, is used for surgical pathology diagnosis.

Algorithms have been developed that perform inference directly on belief networks and influence diagrams through direct manipulation of the graphical structure of the representation. The output of such decision-theoretic inference is the best decision to make, given a patient's preferences, the uncertainties about the world, and the information that is available.

An easy-to-use interface to complex belief-network inference is displayed in Figure 3. Interfaces like this one are employed in a set of decision-support systems developed by a Palo Alto medical-informatics company named Knowledge Industries. The interface offers a straightforward, mouse-directed method for (1) entering symptoms, test results, or other observations (listed in the left panel), (2) for reviewing the specific findings entered for a particular case (middle panel), and (3) for reviewing the output of belief-network computations (right column). The interface also displays recommendations for the next best finding to pursue to minimize uncertainty on the differential diagnosis. We will return to future hands-free interaction for accessing these classes of information.

Human Interface Developments and Directions

Several areas of technical development show great promise for providing better interfaces to medical information, and to the results of computation with automated medical reasoning systems. I shall review several areas of research on advanced human-computer interaction in medicine.

Utility-Directed Displays

There has been recent work at the Section on Medical Informatics at Stanford, and at the Palo Alto Laboratory of the Rockwell Science Center, on the theory and applications of methods for controlling the display of computer-based information in time-critical situations. The techniques developed have application to emergency, surgical, and intensive-care situations, where physicians may have little time to sift through large quantities of potentially relevant information from an automated reasoning system. Techniques have been developed for controlling the tradeoff between the *completeness* and the *complexity* of information displays. That is, a *display reasoner* can custom-tailor the generality or *abstraction* of information presented dynamically so that it contains only the most crucial information, and so that it is at an optimal level of complexity. Typically decision theory is used to select the optimal level of detail of information to display--thus, we use the term *utility-directed* to describe the methodology.

The motivation for utility-directed control of an interface comes from research in the cognitive psychology community, highlighted by the graph in Figure 4, that has demonstrated problems with relaying too much information to decisions makers. As the quantity of relevant information increases, decisions get better, until a maximum is reached. After this point, increasing the amount of relevant information overloads a decision maker. Cognitive psychologists have found that humans cannot consider more than five to nine distinct concepts or "chunks" of information simultaneously. This surprising cognitive limitation was demonstrated in a classic study by Miller (Miller 1956), and has been confirmed repeatedly by experimental psychologists. The capacity of decision makers to consider important influences on a decision may be reduced even further if fast action is demanded in frenetic crisis situations; one cognitive-psychology study showed that people cannot retain and reason simultaneously about more than two concepts in environments filled with distractions (Waugh and Norman 1965).

Intelligent utility-directed display of information has been developed and applied to pathology expert systems (Horvitz, et al. 1986, Horvitz, et al. 1989), to a system that relays information about a physiological model of glucose (Mclaughlin 1987), and to a system for monitoring information from the space shuttle at NASA Mission Control. In some of this work, multiattribute utility has been employed to control the complexity of displays and for queueing and prioritizing the results of diagnostic reasoning (Horvitz 1987).




Figure 4. Cognitive psychologists have demonstrated that interface designers must contend with a fundamental tradeoff between the amount of information displayed to a physician and the degradation of decision quality based on informational overload. (From Horvitz, et al. 1991.)

Figure 5 displays some of the functionality of a system designed under Project Vista, a collaborative research and development effort between the Palo Alto Laboratory and NASA's Johnson Space Center to design computational tools that can manage the complexity of information displayed to human operators in high-stakes, time-critical decision contexts. Although the techniques were developed for controlling the displays of time-pressured engineers at NASA Mission Control in Houston, the methods have applicability in medical decision support. The Vista system monitors the status of five major engine subsystems on the space shuttle. All data from the shuttle is sifted continuously through a belief network (analogous to the sifting of physiological data from a patient). The amount of screen area--and the concomitant detail of information displayed about a particular system--is controlled dynamically by a reasoner that computes the time-criticality of making a response to a possible problem. In use, five areas of the screen expand and contract, jockeying for the attention of the user so as to maximize the utility of an operator's response. The utility-directed control of displays has great promise for becoming an automated interface manager for a large variety of medical decision problems.

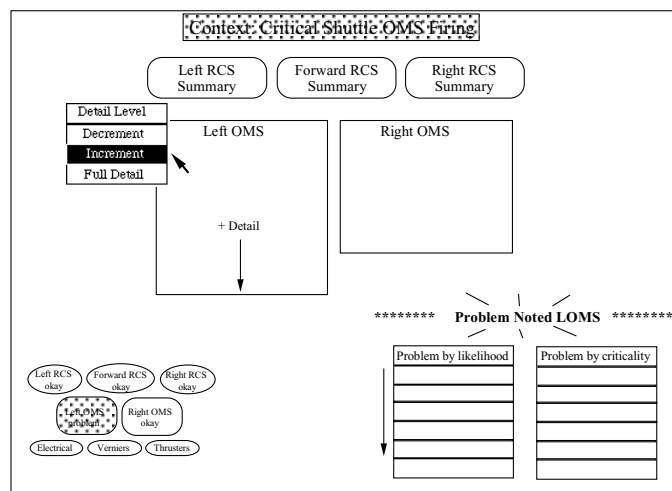


Figure 5. A schematic of the screen generated by a utility-directed display controller implemented for Mission Control at NASA/JSC. At the lower, right corner of the display, a list of possible disorders, computed by a belief network, are listed in terms of the *expected time-criticality* of the problems. The size, and associated

detail of information displayed about five different systems in each of five panels, changes with the time-criticality of the problem. (From Horvitz, et al. 1991.)

User Modeling

Automated explainers should be able to capitalize on knowledge about the user's familiarity with concepts relevant to a computational result. In contrast to short-term memory, long-term memory has been found to be effectively limitless. Thus, previous learning can assist a human with handling complexity. A human that is familiar with a set of ideas or constructs can learn to understand the set as a single concept. This is termed "chunking" in the psychological literature. User modeling refers to techniques for custom-tailoring interactions depending on the expertise or needs of a user. Several techniques have been employed for user modeling. In one approach, employed in the Pathfinder project at Stanford, different sets of costs that represent the difficulty of evaluating patient symptoms accurately have been used, depending on whether the person using the expert system is an expert, a resident, or a medical student.

Hands-Free Systems

The medical-informatics community faces a set of challenging problems in developing information tools that can provide physicians with decision support in areas of medicine dominated by procedures. *Procedure-oriented* tasks include the emergency management of trauma patients and surgery. In such areas, automated reasoning and decision support should not interfere with a physician's hands, and should minimize the obscuration of the physician's view of the patient. We refer to computer systems that can deliver information to physicians as *hands-free decision support*. Key kernels of hands-free computer systems are already under development. A key component of such systems is a head-up display. Such displays allow information to be displayed in a portion of the visual field of a system user. Head-up displays currently under development have the ability to allow a physician to view a computer-generated screen by looking at a predetermined location in the visual field, much as a person with bifocals might look up to gaze at a distant scene.

A prototype system, recently developed by L. Felsenstein, an early pioneer of personal computing systems contains essential components of a hands-free decision-support system. Knowledge Industries and Felsenstein have worked to develop heads-up decision-support systems for use in procedure-oriented problem solving. A user sports a feather-light headset that creates the perception of a computer display at a variable visual distance from the user. A small microphone and earphone allows voice input and output. In addition a small belt-mounted mouse enables a user to move a computer cursor. Future extensions to this system are already in the works. Beyond voice and mouse control of a computer system, gaze-control systems have been created by several research groups. With a gaze-controlled system, the positions of a users pupils and eyelids are monitored, and are used to control the position and activation of a computer cursor. Voice and pupil-driven cursor positioning provides a elegant means for delivering computational decision-support to physicians. Currently, there only about twenty gaze-controlled systems in use in the world.

Hands-free computation systems promise to deliver several classes of decision support to assisting physicians with tasks requiring time-pressured and hands-free access to anatomic and physiologic information. These include (1) the integration and sharing of information about multiple physiologic systems, (2) the access of probabilistic inference to assist with diagnosis and forecasting, and (3) the access of decision-theoretic inference to order priorities, to determine the best information to pursue, and to identify a best therapy plan.

Methods for Information Sharing and Integration

Elegant hands-free interaction technologies can enhance information-sharing and communication among a team of healthcare professionals. The team of players that ascends to care for a newly arrived trauma patient must coordinate procedures and information gathering. To date, much of this coordination is done in a flurry of announcements and commands. Shared hands-free computation systems can allow each member of a trauma team to post a set of findings and to direct a symphony of responses. For example, a team leader can post an ordered set of priorities that can be viewed by all. In addition, all members of the team can share crucial vital signs information. In surgery, hands-free computation systems can allow surgeons to continue to monitor vital signs usually monopolized by the anesthesiologist. The systems can also enable anesthesiologists and nurses to ready themselves for such surgical maneuvers as the administration of local anesthetic or procedures that are likely to lead to great blood loss. Hands-free computation systems can also supply important auxiliary information to surgeons and emergency-room physicians. For example, tomographic images can be displayed to surgeons operating on a cancer patient. Templates or models can be displayed during reconstructive and plastic surgery.

Anatomically-Keyed Displays

Beyond relatively straightforward heads-up displays of computer-generated information, advances at the interface of three-dimensional modeling and medical informatics may someday allow surgeons to inspect a patient, and to perform surgery on a patient, while information generated by real-time tomography or models is overlaid on a patient's anatomy. For example, fragile blood vessels could be highlighted, or the borders of a tumor could be displayed clearly on a patient's anatomy. Computer information can be overlaid on a patient's anatomy by effectively integrating computer-generated images with real-world views of a patient's anatomy. Such *anatomically-keyed* information could enhance the ability of surgeons and other specialists to make use of computer-generated output. Keying information to a patient's anatomy is closely related to ideas developed a decade ago for combining appropriately perspected graphics with live video scenes for professional video productions and broadcasting. Anatomically-keyed displays employ methods for appropriately rotating and perspecting computer-generated overlays. Two basic approaches are the generation of composite images from computer graphics and live video of a patient's anatomy. In another approach, the generation of composite images is created through the optical overlay of computer-based information with special glasses so as to achieve the goal of combining the computer-based data and real-world view of the patient. The latter approach is more sophisticated, requiring machinery to monitor the perspective and gaze of a physician to create the illusion of seeing a composite image.

Telepresence and micropresence

Telepresence, micropresence, and telerobotics hold promise for allowing expert physicians to assist clinicians and surgeons at remote locations. Telepresence allows a physician to have access to a distant location through static or full-perspective views of a patient. High-fidelity telepresence allows a remote location or scene to be inspected from different perspectives via a procedure of moving distant cameras in concert with the head and gaze positions of a local observer. Telerobotics allows the telepresent clinician to interact with a distant patient. Simple forms of telepresence and tele-robotics have already become popular in pathology diagnosis. In telepathology, a surgical pathologist has instant access to a microscope and a slide of a patient's biopsy at a distant site. Telepathology systems that communicate via satellite and over the telephone lines have been developed. Several groups, including teams at NASA, and within the SRI International bioengineering group, have developed interesting demonstration technologies that display the effectiveness of telerobotics for exploring and manipulating objects at a distance. Some telepresence projects demonstrate strides in developing force-feedback techniques, which allow a user to feel the texture, elasticity, or weight of distant objects and structures.

We have been investigating a derivative of telepresence, called *micropresence*, at the Palo Alto Laboratory. Micropresence can enable physicians to explore and to perform procedures on compact, complex regions of a patient's anatomy, while minimizing the extent of surgical incisions. Micropresence involves the positioning of one or more small CCD television cameras and associated camera-control systems in hard-to-reach or compact areas of a patient's anatomy. Such cameras have the ability to image and enlarge complex anatomic regions of interest, as well to identify the exact position of teleoperated microsurgery tools. Micropresence could allow surgeons to explore small or hidden areas from different perspectives, and to perform surgical procedures in these areas as if the regions of interest were expanded greatly, or even made to surround a physician. In essence, a surgeon would be endowed with the ability to explore complex regions, and maneuver tele-operated tools as if he or she were reduced in size, and could step into these regions.

I have touched on several key areas of promise for developing tools that can assist physicians interact more efficiently with computers. We focused, in particular, on technologies that can help physicians make more effective use of computers that offer assistance with decision making. There is great opportunity for enhancing future healthcare delivery by integrating medical-informatics software with evolving human—computer interaction technologies.

Bibliography

(Dagum and Horvitz 1992)

P. Dagum and E. Horvitz. *Computation and Display of Information for Trauma Management and Surgery*. Technical Report, Knowledge Systems Laboratory, Departments of Computer Science and Medicine, Stanford University, March, 1991.

(Heckerman, et al. 1989)

D.E. Heckerman, E.J. Horvitz, and B.N. Nathwani. Update on the Pathfinder project. In *Proceedings of the Thirteenth Symposium on Computer Applications in Medical Care*, Washington, D.C., pages 203-207. IEEE Computer Society Press, Silver Spring, MD, November 1989.

(Henrion, et al. 1991)

M. Henrion, J.S. Breese, and E.J. Horvitz, Decision analysis and expert systems. *AI Magazine* 12(4):64-91.

(Horvitz 1987)

Horvitz, E.J., *A Multiattribute Utility Approach to Inference Understandability and Explanation*. Technical Report KSL-28-87, Knowledge Systems Laboratory, Departments of Computer Science and Medicine, Stanford University, March, 1987.

(Horvitz 1992)

E.J. Horvitz. Automated Reasoning for Biology and Medicine. In R. Fortuner, editor, *Advances in Computer Methods for Systematic Biology: Artificial Intelligence, Databases, and Computer Vision*. Johns Hopkins University Press, Baltimore, MD, 1992.

(Horvitz, et al. 1984)

E.J. Horvitz, D.E. Heckerman, B.N. Nathwani, and L.M. Fagan. Diagnostic Strategies in the Hypothesis-Directed Pathfinder System. In *Proceedings of the First Conference on Artificial Intelligence Applications*, Denver, CO, pages 630-636. IEEE Computer Society, December, 1984.

(Horvitz, et al. 1988)

E.J. Horvitz, J. Breese, and M. Henrion. Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, Special Issue on Uncertainty in Artificial Intelligence, 2:247-302.

(Horvitz, et al. 1989a)

E.J. Horvitz, G.F. Cooper, and D.E. Heckerman. Reflection and action under scarce resources: Theoretical principles and empirical study. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, MI, pages 1121-1127. Morgan Kaufmann, San Mateo, CA, August 1989.

(Horvitz, et al. 1989b)

E.J. Horvitz, D.E. Heckerman, K. Ng, and B.N. Nathwani. Heuristic abstraction in the decision-theoretic Pathfinder system. In *Proceedings of the Thirteenth Symposium on Computer Applications in Medical Care*, Washington, DC, pages 178-182. IEEE Computer Society Press, Silver Spring, MD, 1989.

(Horvitz, et al. 1992)

E. Horvitz, C. Ruokangas, S. Srinivas, and M. Barry. Project Vista: Display of Information for Time-Critical Decisions. In *Proceedings of the Fourth Rockwell Conference on Control Systems and Signal Processing*, Anaheim, CA. Jan, 1992.

(Horvitz and Rutledge 1991)

E.J. Horvitz and G. Rutledge. Time-dependent utility and action under uncertainty. In *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence*, Los Angeles, CA, pages 151-158. Morgan Kaufmann, San Mateo, CA, July 1991.

(Mclaughlin 1987)

J. Mclaughlin. *The Utility-Directed Presentation of Graphical Simulation*. Technical Report KSL-87-59, Knowledge Systems Laboratory, Departments of Computer Science and Medicine, Stanford University. December, 1987.

(Miller 1956)

G.A. Miller. *The Magical Number Seven, Plus or Minus Two*. Psychological Review, 63: 81-97.

(Simon 1973)

H.A. Simon. *The Organization of Complex Systems* in Hierarchy Theory, H.H. Pattee, ed., pages 3-27. G.Brazilier, 1973

(Waugh and Norman 1965)

N.C. Waugh and D.A. Norman. *Primary Memory*, Psychological Review. 72:89-104.