Review ■

A Primer on Aspects of Cognition for Medical Informatics

VIMLA L. PATEL, PHD, DSC, JOSÉ F. AROCHA, PHD, DAVID R. KAUFMAN, PHD

A b s t r a c t   As a multidisciplinary field, medical informatics draws on a range of disciplines, such as computer science, information science, and the social and cognitive sciences. The cognitive sciences can provide important insights into the nature of the processes involved in human–computer interaction and help improve the design of medical information systems by providing insight into the roles that knowledge, memory, and strategies play in a variety of cognitive activities. In this paper, the authors survey literature on aspects of medical cognition and provide a set of claims that they consider to be important in medical informatics.

■ **J Am Med Inform Assoc.** 2001;8:324–343.

Medical informatics is a multidisciplinary field that draws on a range of other disciplines, such as computer science, information science and, more recently, the cognitive and social sciences. Medical informatics is increasingly recognized as being much more than the thin intersection of a science of computing and medical practice.[1] Cognitive science and studies of medical cognition meaningfully inform and shape design, development, and assessment of information systems and decision-support technology.[2] We present a set of exemplars that illustrate how methods and theories from cognitive science can be used to further these objectives.

Technology is increasingly playing a central role in medical practice. As in most workplaces, a gulf exists between technologic artifacts and end users. Bridging this gulf necessitates both changes in the design of artifacts and the development and refinement of human competencies.[3] In addition, a host of issues that extend beyond the use of technology are of concern to investigators in medical informatics. These include the comprehension of medical information by health care personnel, the development of skill and models of competency for complex tasks, and the coordination of knowledge among persons with different backgrounds.

Cognitive science can provide a framework for the analysis and modeling of complex human performance and thereby contribute toward these objectives. It can also provide insight into principles of system usability,[4–7] the process of medical judgment and decision making,[8] the training of physicians and end users, and the study of collaboration in the workplace.[9] This paper offers a primer on the aspects of medical cognition, briefly surveys pertinent literature, and provides a set of claims that can inform researchers in medical informatics.

This primer is not intended to be comprehensive. It addresses a subset of issues that we view as particularly relevant. Areas such as decision making, perceptual diagnosis, and collaborative learning are addressed only briefly here. A survey of medical decision-making research, a vast area of considerable importance, would require an article unto itself. In fact, edited volumes provide extensive reviews of this area.[10,11]

Affiliations of the authors: Columbia University, New York, New York (VLP); McGill University, Montreal, Canada (JFA); and University of California–Berkeley, California (DRK).

Correspondence and reprints: Vimla L. Patel, PhD, Department of Medical Informatics, Vanderbilt Clinic Building, 5th Floor, Columbia University, 622 West 168th Street, New York, NY 10032; e-mail: <patel@dmi.columbia.edu>.

Computer-assisted instruction in medical education is another significant area that intersects with the study of medical cognition. This work is beyond the scope of this paper, and the reader is directed to pertinent reviews.[12,13] Medical discourse and doctor–patient communication is an issue of increasing consequence,[14–16] which also does not receive treatment in this article. Similarly, human factors in medicine and, in particular, the study of medical error[17,18] draw on methods and theories of cognitive science. This impressive body of work is discussed only briefly here.

This paper endeavors to provide a connected thread through foundations, theory, methodology, and research findings in medical cognition. We have emphasized areas of the greatest interest to us, so that work from our laboratory and work related to it are largely represented in this paper.

Medical cognition research looks at the psychological processes underlying performance, focusing on in-depth analysis of the perceptual and cognitive processes that lead to observable behavior. The focus is on understanding the knowledge structures and mental processes brought to bear during cognitive activity (e.g., problem solving and decision making). Although performance varies substantially from person to person, cognitive theories and methods allow us to capitalize on certain structural and processing regularities of the human information processing system, which give strength to generalizations.

Cognitive research strategies tend to vary as a function of theory development. The earlier stages typically necessitate detailed analyses of a few subjects, as researchers endeavor to characterize a new phenomenon. Subsequent research, guided by the initial detailed studies, can employ a larger sample size to test the generality of specific hypotheses. However, in certain cases, the order may be reversed, and large-scale quantitative studies may give rise to in-depth qualitative analyses of a particular phenomenon. For example, studies that document the prevalence of certain decision-making biases maybe followed up by studies that characterize the reasoning processes that contribute to such biases. Similarly, research documenting the failure of many diabetic patients to comply with therapeutic regimens may be followed up by cognitive studies that investigate the beliefs and understandings that sustain such suboptimal practices.

This paper is organized as follows: First, we present a brief history of medical cognition, which serves to situate this field in a broader context and to frame the issues discussed in the rest of the paper. Second, we present a section on foundational issues, in which we discuss structural regularities embodied in human information processing systems; outline an epistemological framework for the study of medical problem solving; and describe cognitive task analysis, a collection of methods designed to characterize the cognitive dimensions of such medical tasks as diagnosis and patient management. Cognitive analysis in knowledge-intensive domains, such as medicine, is predicated on, first, architectural assumptions about the human cognitive system (e.g., memory systems) and, second, systematic analysis of the domain and the nature of the task (e.g., diagnosis).

Finally, the last subsection in the "Foundations" section addresses issues pertaining to the distributed nature of cognition in complex social settings. Distributed cognition, which characterizes how thinking processes occur among groups in specific contexts, has become an increasingly central issue as researchers endeavor to understand cognitive processes in real-world situations.

The second major section involves a review of research findings in medical cognition. This section discusses findings pertaining to the nature of medical expertise, memory and comprehension processes, conceptual understanding, and problem solving and reasoning. In these sections, we touch on important cognitive processes, such as intermediate constructs for facilitating comprehension, the nature of conceptual understanding, and the use of reasoning and problem-solving strategies in diagnostic tasks. We conclude the paper with a discussion of the relevance of investigating cognitive issues for the development and use medical informatics systems.

A Historical Introduction to Medical Cognition

Medical cognition is a subfield of cognitive science devoted to the study of cognitive processes—such as perception, comprehension, reasoning, decision making, and problem solving—in medical tasks. Studies of medical cognition include analyses of performance in "real-world" clinical tasks as well as in experimental tasks. The mysteries of medical cognition have been the subject of discussion for the last several centuries.[19] Understanding the thought processes involved in clinical reasoning in order to promote more effective practices has been the subject of concern for nearly a century.[20] This subdiscipline has also had a close connection to research in medical education.

Medicine is an inordinately complex and multifaceted discipline. Differences in the nature of patient populations and in physicians' educational training and experience, as well as cultural differences, create substantial variation in the nature of medical practices. This provides formidable challenges for those who are interested in the development and implementation of information technologies. It also presents obstacles to researchers interested in understanding medical cognition.

Despite all the idiosyncrasies and particulars, however, we can exploit a remarkable number of regularities to better understand medical cognition. Researchers have made substantial progress in characterizing such regularities in human information processing systems and have similarly made progress in characterizing commonalities in tasks such as diagnostic reasoning and patient management.

Human information processing research has typically focused on the individual. The dual focus on in-depth task analysis and the study of human performance is a central feature of a cognitive science approach. Recently, investigators have made some progress in characterizing dimensions of complex social systems that constitute collaborative performance.[6]

There are two primary approaches to research investigating clinical reasoning in medicine—the decision-analytic approach and the information-processing or problem-solving approach. Decision analysis uses a formal quantitative model of inference and decision-making as the standard of comparison.[11] The typical psychological study of decision-making contrasts the performance of a physician with the mathematical model and generally focuses on the reasoning "fallacies" and biases inherent in human clinical decision-making.[21] The information-processing approach, in turn, focuses on a description of cognitive processes in reasoning tasks, makes use of protocol analysis techniques,[22] and stresses the development of cognitive models of performance. This paper is principally concerned with research related to the information-processing approach.

The systematic investigation of medical expertise began more than 40 years ago with research by Ledley and Lusted[23] into the nature of clinical inquiry. They proposed a two-stage model of clinical reasoning involving a hypothesis generation stage followed by a hypothesis evaluation stage. This latter stage is most amenable to formal decision analytic techniques. Feinstein[24] later proposed an elaborate and ambitious theory of diagnostic reasoning as a logical process. Feinstein was particularly concerned with developing explanatory models of decision making that incorporated physiologic causality, thereby going beyond probabilistic models that considered only the correlations of symptomatology to clinical disorders.

The earliest empirical studies of medical expertise can be traced to the works of Rimoldi[25] and Kleinmuntz,[26] who conducted experimental studies of diagnostic reasoning contrasting students with medical experts in simulated problem-solving tasks. The results emphasized the greater ability of expert physicians to selectively attend to relevant information and narrow the set of diagnostic possibilities (i.e., consider fewer hypotheses).

As cognitive science came into prominence in the early 1970s, spearheaded by the immensely influential work of Newell and Simon on problem solving,[27] research in information-processing psychology accelerated dramatically. Problem solving was conceived of as search in a problem space in which a problem solver is viewed as selecting an option (e.g., hypotheses or inference) or performing an operation (from a space of possible operations) in moving toward a solution or goal state (e.g., diagnosis or treatment plan). This conceptualization had an enormous impact in both cognitive psychology and artificial intelligence research. These constructs enabled researchers both to study patterns of inference making and search strategies in human problem solvers and develop computational models that embody them.

This also led to rapid advances in medical problem-solving research, as exemplified by the seminal work of Elstein et al.[28] They were the first to use experimental methods and theories of cognitive science to investigate clinical competency. Their extensive empirical research led to the development of an elaborated model of hypothetico-deductive reasoning, which proposed that physicians reasoned by first generating and then testing a set of hypotheses to account for clinical data (i.e., reasoning from hypothesis to data). This model of problem solving has had a substantial influence on studies of both medical cognition and medical education.

In the late 1970s and early 1980s, advances into the nature of human expertise were paralleled by developments in medical artificial intelligence, particularly expert-system technology. Although a review of medical artificial intelligence is beyond the scope of this article (but see Shortliffe[29]), we briefly address research that has had a more direct relationship on research in medical cognition. Artificial intelligence in medicine and medical cognition mutually influ-

enced each other in a number of ways, including providing a basis for developing formal models of competence in problem-solving tasks; elucidating the structure of medical knowledge and providing important epistemological distinctions; and characterizing productive and less-productive lines of reasoning in diagnostic and therapeutic tasks.

Gorry[30] conducted a series of studies examining the ways in which a computational model of medical problem solving compared with the actual problem-solving behavior of physicians. This analysis provided a basis for characterizing a sequential process of medical decision making, one that differs in important respects from early diagnostic computational systems based on Bayes' theorem. Pauker et al.[31] capitalized on some of the insights of Gorry's earlier work and developed Present Illness Program (PIP), a program designed to take the history of a patient with edema. Several questions guiding this research, including the nature and organization of expert knowledge, were of central concern to both developers of medical expert systems and researchers in medical cognition. The development and refinement of the program was based partly on studies of clinical problem solving.

Medical expert consultation systems such as Internist,[32] MYCIN,[33] and ABEL[34] seeded ideas about knowledge-based reasoning strategies across a range of cognitive tasks. MYCIN, in particular, had a substantial influence on cognitive science. It contributed several advances (e.g., representing reasoning under uncertainty) in the use of production systems as a representation scheme in a complex knowledge-based domain. The inference engine, EMYCIN, was widely used in the development of other expert systems and as a cognitive modeling tool. MYCIN also served to highlight the cognitive dimensions of medical explanation as distinct from medical problem solving.

The work of Clancey and Letsinger[35,36] in the context of developing first GUIDON and then NEOMYCIN was particularly influential in the evolution of models of medical cognition. Clancey endeavored to reconfigure MYCIN to employ the system to teach medical students about meningitis and related disorders. NEOMYCIN was based on a more psychologically plausible model of medical diagnosis. This model differentiated data-directed and hypothesis-directed reasoning and separated control knowledge from the facts it operates on. In particular, it distinguished among findings, hypotheses, evidence (finding–hypotheses links), justifications (why a finding–

hypotheses link is true), structure (how findings and hypotheses are related among themselves), and strategy (why a finding request comes to mind).

Feltovich et al.,[37] drawing on models of knowledge representation from medical artificial intelligence, characterized fine-grained differences in knowledge organization between subjects of different levels of expertise in the domain of pediatric cardiology. These differences accounted for subjects' inferences about diagnostic cues and evaluation of competing hypotheses. Patel and Groen,[38] incorporating distinctions introduced by Clancey, studied the knowledge-based solution strategies of expert cardiologists as evidenced by their pathophysiologic explanations of a complex clinical problem. The results indicated that subjects who accurately diagnosed the problem, employed a forward-oriented reasoning strategy, using patient data to lead to a complete diagnosis (i.e., reasoning from data to hypothesis). This is in contrast to subjects who misdiagnosed or partly diagnosed the patient problem. They tended to use a backward reasoning strategy. The results of this study presented a challenge to the hypothetico-deductive model of reasoning espoused by Elstein et al.,[28] which did not differentiate expert from non-expert reasoning strategies.

Perceptual diagnosis has also been an active area of inquiry. Studies have examined the abilities of subjects at several levels of expertise to diagnose skin lesions presented on a slide.[39] The results revealed a monotonic increase in accuracy as a function of expertise. In a classification task, novices categorized lesions by their surface features (e.g., scaly lesions), intermediates grouped the slides according to diagnosis, and expert dermatologists organized the slides according to superordinate categories such as viral infections, which reflected the underlying pathophysiologic structure.

The ability to abstract the underlying principles of a problem is considered one of the hallmarks of expertise, both in medical problem solving and in other domains.[40] Lesgold et al.[41] investigated the abilities of radiologists at different levels of expertise to interpret chest x-ray pictures. The results revealed that the experts were able to rapidly invoke the appropriate schema and initially detect a general pattern of disease, which resulted in a gross anatomic localization and served to constrain the possible interpretations. Novices experienced greater difficulty focusing on the important structures and were more likely to maintain inappropriate interpretations despite discrepant findings in the patient's history.

Foundations

Human Information Processors: Memory and Knowledge Systems

Claim 1: The regularities of the human information processing system enable us to develop reasonably precise expectations and to draw broad generalization concerning cognitive performance.*

Cognitive science serves as a basic science to, and provides a framework for, the analysis and modeling of complex human performance. A computational theory of mind provides the fundamental underpinning for most contemporary theories of cognitive science. The basic premise is that much human cognition can be characterized as a series of operations or computations on mental representations. Mental representations are internal cognitive states that have a certain correspondence with the external world. For example, they may reflect a belief about a patient after the patient's somewhat yellowish skin pallor is noticed. These mental representations are likely to elicit further inferences about the patient's underlying condition and may direct the physicians' information-gathering strategies. This is a simple illustration of an operation on a representation.

Individuals differ substantially in terms of their knowledge, experiences, and endowed capabilities. Nevertheless, we can characterize certain regularities of the human information-processing system. These can be either structural regularities—such as the existence of and the relations between perceptual, attentional, and memory systems and memory capacity limitations—or processing regularities, such as processing speed, selective attention, or problem-solving strategies. Cognitive systems are characterized functionally in terms of the capabilities they enable (e.g., focused attention on selective visual features), the way they constrain human cognitive performance (e.g. limitations on memory), and their development during the life span.[42] An extended discussion of this complex subject matter is beyond the scope of this paper. The reader is referred to other sources for further detail.[43–46] In this paper, we briefly address some important distinctions in the study of human memory and knowledge.

Human memory is typically divided into at least two structures, long-term memory and short-term, or working, memory.[47] Long-term memory (LTM) can be thought of as a repository of all knowledge, whereas working memory (WM) refers to the resources needed to maintain information active during cognitive activity (e.g., problem solving). The information maintained in working memory includes stimuli from the environment (e.g., words on a page) and knowledge activated from long-term memory. In theory, LTM is infinite, whereas WM is limited to five to ten "chunks" of information (that is, variable amounts of information that can be recalled at one time).[48] Problems impose a varying cognitive load on working memory. For example, maintaining a seven-digit phone number in WM is not very difficult. However, attempting to maintain a phone number while engaging in conversation is nearly impossible for most people.

The kinds of knowledge that reside in LTM can be characterized in several ways. First, conceptual knowledge can be distinguished from procedural knowledge. *Conceptual knowledge* refers to one's understanding of domain specific concepts. It supports explanations and may result in appropriate actions. *Procedural knowledge* relates to how to perform various activities. Cognitive and behavioral skills involve procedural knowledge. Numerous technical and other skills in medical contexts require the acquisition of procedural knowledge.

Conceptual knowledge and procedural knowledge are acquired through different learning mechanisms. Conceptual knowledge is acquired through mindful engagement with materials in a range of contexts. Procedural knowledge is developed as a function of deliberate practice, which results in a learning process known as knowledge compilation. However, the development of skills may involve a transition from a declarative or interpretive stage toward increasingly proceduralized stages. For example, in using an electronic medical record (EMR) system during consultation with a patient, a less experienced user will need to attend carefully to every action and input, whereas a more experienced user can more effortlessly interview a patient and simultaneously record patient data.[5,49,50]

Procedural knowledge is rather limited in its generality. A skilled user of one type of EMR system, for instance, will perform less well when using a different system. The extent of this decrease is partly a function of the similarity and differences between the two systems and the kinds of component processes that they employ. To take a simple example, pen-based or mouse-based point-and-click EMR systems engage different motor and perceptual processes than do keyboard-based systems.

---

*This paper is built around a set of claims. We use the term "claims" to denote hypotheses that have been subjected to repeated empirical testing. Claims have less epistemic force than "assertions" or "principles," in that they are subject to revision or refutation.

In addition, factual knowledge can be differentiated from conceptual knowledge. Factual knowledge involves merely knowing a fact or set of facts (e.g., risk factors for heart disease) without any in-depth understanding. Facts are current truth and may become rapidly out of date. The acquisition of factual knowledge alone is not likely to lead to any increase in understanding or behavioral change.[51]

The acquisition of conceptual knowledge involves the integration of new information with prior knowledge and necessitates a deeper level of understanding. For example, risk factors may be associated in the physician's mind with biochemical mechanisms and typical patient manifestations. On the other hand, a clinician who has merely factual knowledge about the empirical relationships between common manifestations and disease may not appreciate atypical presentations or attend to significant findings in the patient's history (e.g., drug allergies). Factual knowledge is inherently more brittle than conceptual knowledge, and this brittleness is most acutely observed in unfamiliar situations.

The regularities of the human information-processing system allow us to draw broad generalizations about how individuals will respond to a cognitive task. For example, a teaching strategy that suggests to medical students that they pursue many diagnostic hypotheses simultaneously is likely to lead to suboptimal learning, because this approach to diagnosis imposes a significant load on memory. It is likely to divide students' attentional and cognitive resources and render learning more difficult.

We can draw analogous conclusions about how different people will design a computer system. For instance, studies of computer interface design have shown that some systems create a substantial burden on working memory,[52] since the user has to process more information than his or her memory capacity allows. Similarly, a computer-based system that transforms a normally routine task, such as recording a laboratory result, into a complex and time-consuming one is likely to increase a user's frustration and diminish his or her interest in continuing to use the system.

Fortunately, memory limitations can be overcome by long and deliberate practice over extended periods of time.[53,54] As people become increasingly familiar with some knowledge domain or task, they acquire a great deal of specific knowledge. In addition, people learn to use their LTM as support for their WM, so that they do not need to remember a large number of isolated pieces of information. Instead, information maintained in WM is strongly linked to other information in LTM by retrieval structures. In a problem-solving situation, an expert can use these retrieval structures to provide selective and rapid access to long-term memory.

For example, when a physician is performing a procedure, the information most easily remembered serves to cue other associated information in LTM, which can then be retrieved without difficulty. Retrieval structures also account for the ability of chess masters to play simultaneously against several opponents at once,[55] since they do not need to memorize the position of each chess piece on the board but can recognize patterns of positions in chunks. Similarly, an expert physician may be able to diagnose a patient's problem by recognizing whole patterns of clinical findings and associating them with particular disease patterns stored in LTM. This precludes the need to process each clinical finding separately.

The extent to which people commit to deliberate practice in a domain can result in significant individual differences in performance and knowledge, which critically affect information processing and the kinds of representations that they are likely to form. The nature of differences in knowledge and performance is a central issue in expert–novice research and is discussed in greater detail below.

For the present purposes, consider a simple example. A man accompanies a friend to a jazz club for the first time. He is unfamiliar with jazz and experiences the music as a series of fragments of familiar melodies alternating with harsh dissonant sounds. His friend, who is rather knowledgeable about jazz and its history, forms a very different set of representations. The more expert listener has a different set of expectations, can selectively attend to different aspects of the music, and can perceive coherence where the novice hears not much more than the sum of individual instruments. Similarly, an expert radiologist can selectively attend to critical features and perceive a lung tumor of a certain kind in a complex and noisy x-ray picture. However, a novice is likely to have difficulty disambiguating shadows from tissue and discerning normal from pathologic anatomic details. Knowledge enables us to process information in very different ways and selectively attend to what is important while filtering out what is not.

An Epistemological Framework for Medical Comprehension and Problem Solving

Claim 2: Medical knowledge can be construed as a hierarchically organized conceptual system that serves to partition problems into manageable clusters of information.
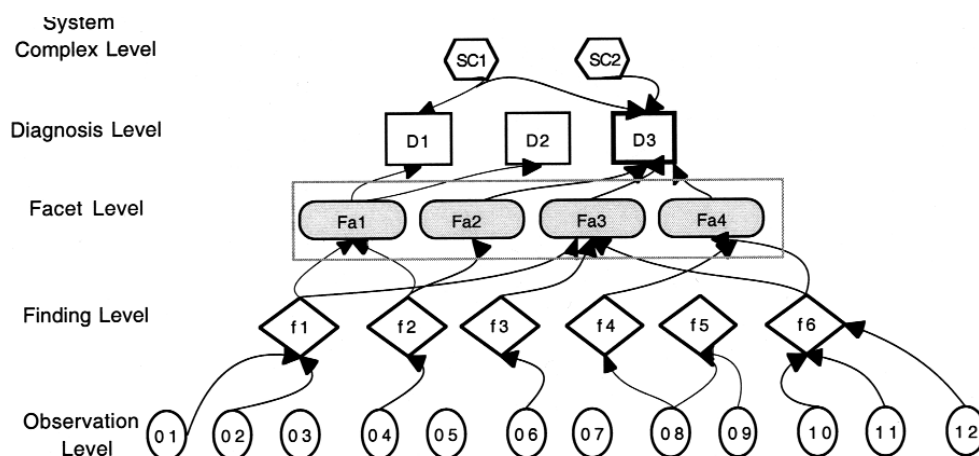
Figure 1 Epistemological framework representing the structure of medical knowledge for problem solving.

What are the central units of medical knowledge? Clearly, this question has no single answer, but the issue is of considerable importance in research in medical informatics as well as in the study of medical cognition. Much research has been conducted in medical artificial intelligence with the aim of developing medical ontologies for use in expert systems.[56,57] In a similar vein, Evans and Gadd[58] proposed a framework that serves to characterize the knowledge used for medical understanding and problem solving and also for differentiating the levels at which medical knowledge may be organized. This framework represents a formalization of medical knowledge as realized in textbooks and journals and can be used to provide insight into the organization of clinical practitioners' knowledge.[†]

The framework consists of a hierarchic structure of concepts formed by observations at the lowest level, followed by findings, facets, and, diagnoses, as illustrated in Figure 1. *Observations* are units of information that are recognized as potentially relevant in the problem-solving context. However, they do not constitute clinically useful facts. *Findings* comprise observations that have potential clinical significance. Establishing a finding reflects a decision made by a physician that an array of data contains a significant cue or cues that need to be taken into account. *Facets* consist of clusters of findings that indicate an underlying medical problem or class of problems. They reflect general descriptions of pathologic conditions, such as aortic insufficiency or endocrine disorder. Facets resemble the kinds of constructs used by researchers in medical artificial intelligence to describe the partitioning of a prob-

lem space.[59,60] They are interim hypotheses that serve to divide the information in the problem into sets of manageable sub-problems and to suggest possible solutions. Facets also vary in terms of their levels of abstraction. *Diagnosis* is the level of classification that subsumes and explains all levels beneath it. Finally, the *systems level* consists of information that serves to contextualize a particular problem, such as the ethnic background of a patient.

An empirically tested epistemological framework such as this allows us to characterize differential organization of information in a range of tasks, from writing clinical case summaries to performing diagnostic reasoning. To this end, the framework has been used as a basis for building reference models for medical knowledge, for coding inferences in studies of medical text comprehension,[61] and for characterizing clinical reasoning[62] and doctor–patient interaction.[63]

These reference models, which are idealizations or benchmarks of accurate reasoning, can then be mapped to physicians' or medical students' representations of clinical information. For example, we can compare the findings elicited by physicians who correctly diagnosed a problem with those of physicians who misdiagnosed the problem. We can also consider how physicians cluster findings to form facets and subsequently generate and test diagnostic hypotheses. Research results have shown[62,64] that physicians understand and reason about clinical information at different levels, depending on their levels of expertise. Knowledge of these differences can be useful for the design of decision-support or training systems in which information (e.g., reminders, help) could be presented at multiple levels of the aggregation. In this way, information can be matched to the level of medical knowledge a user is more likely to understand.[65]

[†]We do not claim that this hierarchic representation is the only way in which physicians organize medical knowledge. In fact, physicians have access to clinical exemplars and "deep" biomedical knowledge, which provide alternative forms of representation.

Cognitive Task Analysis

> *Claim 3:* Cognitive task analyses allow us to define the demand characteristics of a particular task and to focus attention on the pertinent problem dimensions.

The study of human performance is predicated on an analysis of both the information-processing demands of a task and the kinds of domain-specific knowledge required to perform it. This analysis is often referred to as cognitive task analysis. The principles and methods that inform this approach can be applied to a wide range of tasks, from the analysis of written medical texts to the investigation of EMR systems. We can characterize sets of generic tasks that necessitate similar cognitive demands. For example, clinical tasks in medicine include diagnostic reasoning, therapeutic reasoning, and patient monitoring and management. Each of these tasks, which may vary substantially in different domains of medicine and situations, has a common underlying structure that requires similar kinds of reasoning and patterns of inference.

The notion of a generic cognitive task originates in the area of medical artificial intelligence.[66–68] It is predicated on the notion that intelligent agents confront similar computational problems, and the information-processing strategies used to solve the problems are essentially similar. Each strategy employs a characteristic inference procedure suitable to the task.

Josephson and Josephson[67] present an analysis of how the computationally complex problem of diagnosis can be decomposed into a set of elementary generic tasks. Diagnostic problem solving is viewed as a task in which the goal is to explain a set of observations of a system. The explanation takes the form of malfunction hypotheses, which may account for observed manifestations. Gott and Lesgold[69] employ cognitive task analyses in complex technical work settings (e.g., avionics). The goals of their analyses are oriented principally toward assessment of performance and instruction. Their analyses yield performance models of technical expertise sufficiently detailed to be used as the benchmark performance targeted by instructional systems and can also elucidate potential impasses or barriers in the process of skill acquisition.

In a cognitive task analysis, some tasks be decomposed into constituent parts. We conducted a set of studies in which physicians use EMR systems while concurrently interviewing patients.[49,50] Employing the EMR system to enter patient data engages the component processes of information gathering, diagnostic reasoning,[‡] and problem representation. Each of these processes can be decomposed into set of goals and action sequences. The structure of the information gathering focuses on the kinds of information obtained during the clinical history, physical examination, and laboratory investigations. For example, if the goal is to conduct a comprehensive clinical interview for a patient complaining of blurred vision and weight loss, a set of subgoals would include characterizing the chief complaints and history of present illness; conducting a review of systems with particular attention to neurologic, ophthalmologic, endocrine, and other systems; and obtaining a family history, focusing on possible histories of cancer, diabetes, and thyroid problems.

The diagnostic reasoning process includes translation of the "free speech" of the patient or family into "standard tokens" (i.e., a finding) understood and used by medical professionals and used in clinical literature. For example, a patient may provide the information that "he has been feeling tired lately and seems to get out of breath when he climbs the stairs or shovels snow in his driveway." These observations correspond to the finding "exertional dyspnea." Similarly, a particular laboratory result needs to be interpreted in a particular context before a finding can be inferred. Diagnostic reasoning also includes *data evaluation* processes (which can involve characterizing the importance of a finding or qualifying it, for example, in terms of its severity) and the definition of *hypothesis formation and evaluation* goals that involve the development, assessment, and confirmation or rejection of diagnostic hypotheses.

The problem-representation task involves translating clinical findings into a form that is consistent with the kinds of inputs acceptable to the system. It varies considerably depending on the nature of the EMR. For example, highly structured EMR systems necessitate that the user map a clinical finding according to dictates of the underlying controlled medical vocabulary and the screen display. In contrast, a system that is more reliant on free-text entry will invoke a very different set of cognitive and behavioral responses.[49]

Each component task can be further partitioned into increasingly fine-grained goal-action sequences. It is important to note that a cognitive task analysis can specify the abstract structure of a task. In practice, performance of a task may differ for any number of reasons. In addition, clinical situations are fluid rather than static, and goal-action responses may change dynamically as a situation dictates (e.g., with a sudden change in priorities).

[‡]This presupposes that diagnosis is the purpose of the interview. We recognize that EMRs are used for a wide range of tasks.

Distributing Cognition Across Social Agents and Material Artifacts

> *Claim 4:* Routine human cognition is embedded in social practices and is transformed via the use of technology and other cultural artifacts.

In recent years, the conventional information-processing approach has come under criticism for its narrow focus on the rational or cognitive processes of the solitary individual. One of the most compelling proposals has to do with a shift from viewing cognition as a property of the individual to viewing it as distributed across groups, cultures, and artifacts. This claim has significant implications, not only for cognitive theory and research but also for instructional and technologic design.

Cole and Engeström[70] suggest that "a natural unit of analysis for the study of human behavior is activity systems, historically conditioned systems of relations among individuals and their proximal, culturally organized environments." In this view, the individual, groups of individuals, and artifacts can be construed as a single indivisible cognitive agent.[71–73] Lave[74] and others have claimed that activities are so context-bound and that the processes involved in an activity are so varied from one context to another that the distinction between the individual's cognition, the activity, and the context become untenable. According to this proposal, the only meaningful level of analysis is in terms of a person-acting-in-a-setting, with the individual contributors forming an integral part of the composite. These issues have been addressed in detail in relation to cognition and education[75,76] and have been examined in the context of medical informatics.[77,78]

The situated critique is a compelling one, and it serves to highlight the limitations of an approach that focuses merely on the solitary individual. Much everyday cognition is embedded in social practices that involve other participants and artifacts. However, we tend to view the approaches as complementary. The information-processing perspective offers a powerful set of methodological and theoretic tools to understand cognition, whereas the distributed approach (which embodies several distinct traditions, such as cultural-historical activity theory) helps us understand how social entities jointly make (or distribute) decisions and cognitive resources.

The distributed approach also provides a basis for understanding how communication is grounded (how meaning is jointly negotiated) across geographic distances, how groups jointly learn, and the ways in which organizational entities are constituted to produce (and sometimes impede) work. Suchman,[79] Vicente,[80] and others have provided us with penetrating analyses for understanding how cognition, activity, social relationships, and learning are mutually constituted in a workplace setting. Hutchins,[73,81] in particular, provides theoretic and methodological insights into how researchers can characterize the ways in which a social entity—such as ship navigators, airplane pilots, or surgical teams—"fuses" with technologies and artifacts (e.g., patient charts, monitoring devices, graphic representations) to form a single, indivisible information-processing system.

We have previously drawn on these new approaches to characterize decision making in critical care settings[82] and the process of computer-mediated collaborative design in medical informatics.[6] From a technologic perspective, the enormous impact of artifacts on human cognition is being recognized, which can lead to a greater understanding of the information-processing roles played by artifacts and how they interact with the information processing activities of their users.[83,84] From a social perspective, this approach signals a need to develop new methodologies and refine older ones, to characterize the division of labor and the distribution of effort and cognitive resources that constitute most cognitive activities. This perspective maintains continuity with symbolic information-processing theory, thereby building a rich repository of research findings, methods and theories. The discipline of applied cognitive science is beginning to develop integrated approaches for combining in-depth examination of individual cognition with characterization of the broader social and material context.

We can characterize human performance in social settings at various levels of analysis, from the study of the technology itself (e.g., a new medical information system) to increasingly larger arenas such as medical teams, hospital sites, and the larger community. Each level of analysis addresses a particular aspect of a problem.

A given study may involve two or more such layers. Consider the study of a new medical information technology. The most basic research can be carried out on the system itself. In fact, usability-inspection techniques, such as the cognitive walk-through, [85] analyze the ease and efficacy with which a given system can be used to perform particular tasks. The next layer includes a single person working with a system. Typical usability testing involves a representative user (e.g., a nurse, physician, or other health care worker) performing a task that the system is designed to support.[5]

The first two layers have constituted the bulk of research in human–computer interaction, although the third one, involving collaboration among multiple users, has increasingly been a subject of focal concern. The outer layers introduce additional kinds of users in diverse settings and substantial complexity. The third layer might represent a hospital clinic in which many persons jointly use the system. Subsequent layers could include the entire hospital, the community served by this establishment, and regional health care organizations. Analysis of each layer requires different methodologies, drawing on a range of social and behavioral-science disciplines. The levels of analysis constrain one another, and the composite provides a much richer picture of the systemic processes (e.g., medical decision making in a complex real-world setting).

Research

Individual Differences and the Development of Expertise

> *Claim 5:* The development of expertise is characterized by a systematic but "non-monotonic" growth in knowledge and performance, in which the progress from a novice to an expert does not necessarily increase with training or time on task.

Most research in medical cognition has employed the expert–novice approach, which involves the comparison of people at different levels of training, experience, and mastery in the performance of a variety of health-related tasks. Two major goals of this research have been to understand what distinguishes outstanding individuals in a domain from less outstanding individuals[86] and to characterize the development of expertise. This approach originated with the pioneering research of deGroot[87] in the domain of chess, from which it extended to investigations of expertise in a range of content domains, including physics[40,88] music,[89] and medicine.[64] This research has shown that, on average, the achievement of expert levels of performance in any domain requires about ten years of full-time experience. An "expert" is someone who has achieved a high level of proficiency, as indicated by various measures, such as ELO ratings in chess, world rankings in various athletic endeavors, and certification by a sanctioned licensing body, as in medical subspecialties.

In medicine, the expert–novice paradigm has contributed to our understanding of the nature of medical expertise and skilled clinical performance.[41,64] Expert physicians have extensive general knowledge of medicine (acquired through medical school and residency training) and deep, detailed knowledge of their relatively narrow areas of specialization (acquired partly from training and clinical experience). We can distinguish between *specific* (e.g., cardiology) and *generic* (e.g., general medicine) expertise. Every experienced physician has acquired common wisdom and medical knowledge as well as certain mastery in the application of medical skills; this constitutes generic expertise. It is hypothesized that a person may possess both kinds of expertise or only generic expertise. Early medical training, through medical school and internship, largely emphasizes the acquisition of generic expertise.

When physicians enter a residency training program, they begin to specialize, thereby acquiring specific expertise as well as continuing to develop generic expertise. A classification of different levels of expertise can be as follows[64]: A *beginner* is a person who has only everyday, lay knowledge of a domain; an example is a typical patient. A *novice* is someone who has begun to acquire the prerequisite knowledge assumed in the domain, such as a medical student. Novices have a basic familiarity with the core concepts, the language and, to a lesser extent, the culture of medicine. An *intermediate* is above the beginner level but below the subexpert level and is typically a senior medical student or a junior resident. A *subexpert* (e.g., a specialist solving a clinical problem outside his or her domain of expertise) possesses generic knowledge but lacks specialized knowledge of a medical subdomain. Finally, an *expert* (e.g., a cardiologist or an experienced intensive care nurse) has specialized knowledge of the subdomain.

The development of expertise typically follows a somewhat unusual trajectory. It is often assumed that the novice becomes an expert by a steady, gradual accumulation of knowledge and fine-tuning of skills. That is, as a person becomes more familiar with a domain, his or her level of performance (e.g., accuracy and quality) gradually increases.

Research has questioned this assumption, however.[41,62,90] Cross-sectional studies of experts, intermediates, and novices have shown that, on some tasks, people at intermediate levels of expertise may perform more poorly than those at lower levels of expertise, a phenomenon known as the "intermediate effect."[90] When novice–intermediate–expert data are plotted, they resemble the pattern shown in Figure 2, in which the performance of intermediate subjects (those who are on their way to becoming proficient in a domain but have not reached the level of experts) declines to a
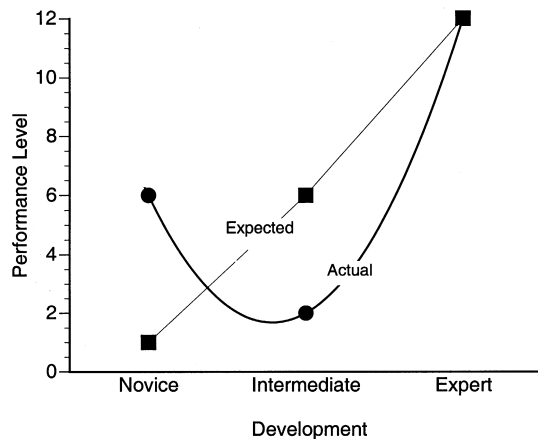
Figure 2 Idealized representation of the intermediate effect, a finding reported in many studies of expertise, learning, and cognitive development. The straight line gives a commonly assumed representation of performance development by level of expertise. The curved line represents the actual development from novice to expert. The Y-axis may represent any of a number of performance variables, such as the number of errors made, number of irrelevant concepts recalled, number of conceptual elaborations, or number of extraneous hypotheses generated in a variety of tasks.

level below that of novices. This progress from novice to expert is said to be non-monotonic, because it does not necessarily increase with training.

We have found evidence of the intermediate effect in a variety of tasks, including recall of patient data,[90] explanation of clinical problems,[91] doctor–patient communication,[63] explanation of laboratory data,[64,92] and the generation of diagnostic hypotheses.[93] In all these tasks, intermediates appear to perform below the level of the novices. Similar results have been observed by other investigators in studies of recall of clinical case information,[94] diagnostic reasoning in pediatrics,[95] and perceptual diagnosis.[41] Furthermore, a phenomenon similar to the intermediate effect has been found in other domains, such as chess,[96] and in the learning[97] and developmental literature.[98]

This literature suggests that human development and learning does not necessarily consist of the gradually increasing accumulation of knowledge and skills. Rather, it is characterized by the arduous process of continually learning, re-learning, and exercising new knowledge, punctuated by periods of apparent decrease in mastery and declines in performance. Given the ubiquity of this phenomenon, we can argue that such decline may be a necessary part of learning.

The investigation and use of information technologies in support of training and education have a long history. However, few studies have been based on an

understanding of how people learn and the cognitive processes that underlie learning and performance. Cognitive science has identified learning phenomena, such as the intermediate effect, that have to be taken into account in the design of instruction. This is particularly important for the design and development of training systems. To improve usability and effectiveness, such design should be informed by an understanding of the way people learn and the theories explaining such learning. This can be achieved by developing systems that aid the intermediate student in organizing knowledge in more efficient ways (e.g., by introducing structure to what seems to the intermediate student a set of unrelated findings).

Medical Comprehension and Memory

Claim 6: Comprehension of information involves the development of schemata, which serve to filter relevant information selectively, thereby circumventing limitations of memory.

Comprehension refers to cognitive processes associated with understanding or deriving meaning from text, conversation,[99] and other informational resources.[100] It involves the processes that people use when trying to make sense of a piece of text, such as a sentence, a book, or a verbal utterance. It also involves the final product of such processes, that is, the mental representation that people have of a text (what they have understood). Comprehension often precedes problem solving and decision making, but it is also dependent on perceptual processes that focus attention, the availability of relevant knowledge, and the ability to deploy knowledge in a given context. In fact, some of the important differences in medical problem solving and decision making arise from differences in comprehension.[62]

The study of comprehension is intimately linked to the study of memory. Indeed, memory tasks are often used to investigate people's comprehension of information. In cognitive science, it is assumed that prior knowledge is stored in long-term memory (LTM). Some of this information is structured as schemata, which are mental representations of typical things (e.g., diseases) and events (e.g., episodes of illness). When a person interprets information, the schema serves as a "filter" for distinguishing relevant and irrelevant information.

Schemas can be considered generic knowledge structures that contain slots for particular kinds of findings (data). For instance, a schema for myocardial infarction most likely contains the findings "chest pain," "sweating," and "shortness of breath" but not the finding

"goiter," which is part of the schema for thyroid disease. To characterize the nature of such schemata and to identify differences between novice and expert comprehension of clinical information, the epistemological framework of medical knowledge described in the previous section can be used. In particular, this framework allows the identification of the level of knowledge that is generated by novices, intermediates, and experts when they interpret clinical problems.

A function of schemata is to provide a "filtering" mechanism to experts, allowing them to selectively attend to significant information and discard irrelevant clinical information.[64] Intermediates are often found to display a great deal of knowledge (sometimes more than an expert). However, they engage in excessive and irrelevant elaboration or search for information when attempting to understand or explain a clinical problem. In contrast, experts' knowledge is finely tuned to the performance of various tasks, and they can readily filter out irrelevant information using their schemata (characterized in terms of findings, facets, diagnoses, and global complexes). This tuning is a function of the organization of expert knowledge in a functional manner, that is, in a way that is readily available in specific situations.[62]

Intermediate Constructs and Memory

Claim 7: An expert's knowledge enables efficient processing of complex and superficially diverse information. Intermediate constructs are dynamically generated knowledge clusters that serve to partition problems and enable such efficient processing by reducing memory load.

Experts process information at a level of abstraction that is most efficient and reduces the burden on memory. Through years of experience, they have learned to conceptualize medical information (e.g., clinical findings from a patient) in terms of constructs that correspond to facets (intermediate between the concrete level of particular signs and symptoms and the more abstract nature of diagnoses) in the epistemological framework presented previously. In contrast, less experienced physicians tend to process medical information at a more detailed level.[62,101] Facets have a broad explanatory power and provide coherence to a cluster of findings. They also allow experts to integrate biomedical and clinical knowledge in a manner that fosters efficient processing, leading more directly to the correct diagnosis.

This intermediate level of aggregation used by the expert when interpreting clinical information has some advantages. One such advantage is that interme-diate constructs naturally integrate theoretically based biomedical knowledge with practical clinical knowledge, since they fall between two levels of granularity (diagnoses and findings). General biomedical principles are too abstract and are therefore of limited applicability in practice, whereas clinical exemplars, such as analogies or reminders of prior cases, may be too context-specific to be used in problems.

To illustrate, we offer a diagram that shows inferences made by an expert cardiologist and a subexpert (endocrinologist) in response to the finding of "shortness of breath" in a patient (Figure 3). The expert generates three high-level hypotheses that constrain the type of problem the patient may have (heart problem, respiratory problem, and severe anemia), while the subexpert generates a high-level hypothesis (cardiac disease) but also generates the more specific hypothesis of emphysema." By generating these three facets, the expert reduces the burden on memory by filtering out types of diseases involving other organ systems. As new patient information is presented, the expert selects among the three facets until finally choosing one.

In the next segment (problem segment 3), the expert generates the correct problem type (pericardial problem) by focusing on only one finding (breathless lying down), while the subexpert is still generating more specific hypotheses that lead him away from the correct diagnosis. By problem segment 15, the expert has already rejected previously generated hypotheses and is checking the hypothesis of pericardial disease by generating findings that would be consistent with this type of problem (e.g., ascites). In contrast, the subexpert generates still other hypotheses. The result of the expert's approach is that he or she narrows down to one type of problem very early in the case and then tests his or her facet-hypothesis against the patient data. In contrast, the subexpert keeps generating specific hypotheses, which tend to lead him astray.[62,101]

In summary, expertise in medicine and in other domains involves the construction of clusters of concepts that have great explanatory power (facets) and serve to partition a problem space. These concepts are intermediate between data, such as patient findings, and hypotheses, such as the final diagnosis given to a clinical problem. The use of such intermediate concepts allows experts to remember information in a concise manner. The potential power of intermediate models has also been recognized by the artificial intelligence in medicine community.[102] Diagnostic systems based exclusively on shallow representations (e.g.,
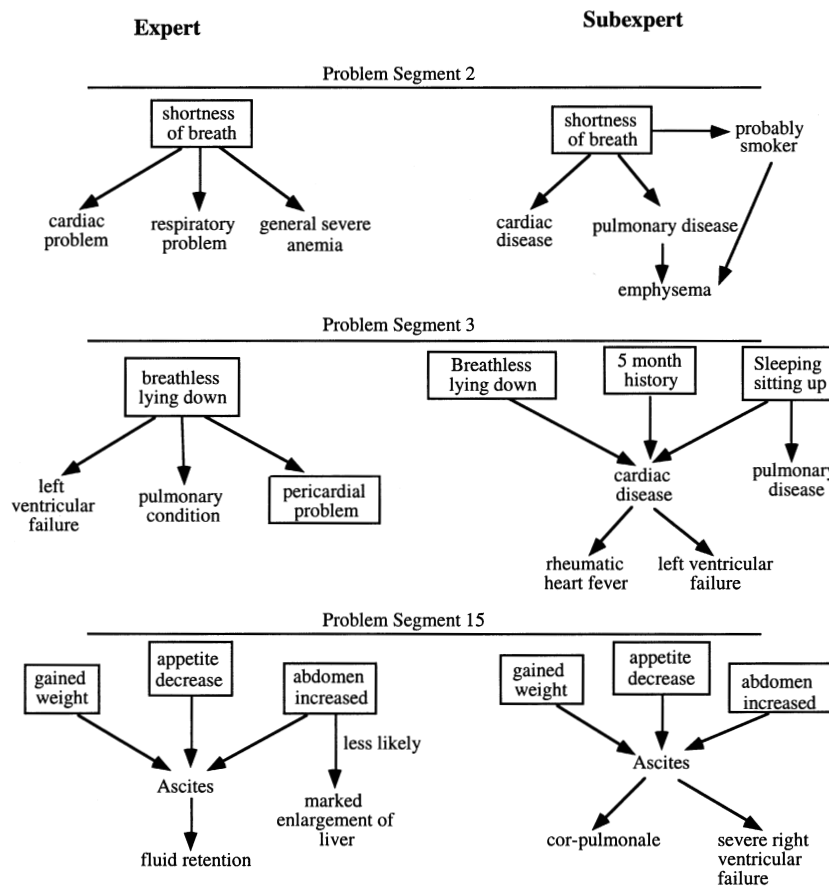
Figure 3 Comparison between the concepts generated by a cardiologist (expert) and an endocrinologist (subexpert) on three pieces of clinical information (problem segments) specific to a cardiology task.

patterns of findings associated with diseases) are inherently brittle, whereas systems based on "deeper" qualitative models (e.g., a detailed model of the cardiovascular system) are computationally expensive. Coiera[102] proposed that qualitative representations of medical histories of intermediate depth offer a distinct computational advantage. Models of intermediate levels of abstraction can also bridge the gap between concrete exemplars such as clinical cases and abstract scientific principles such as the Frank–Starling Law of the Heart, an important principle governing the regulation of cardiac output. They enable a clinician to predict and explain various complex patient manifestations rapidly and efficiently.

The utility of intermediate constructs has been demonstrated in other domains. For example, White[103] and colleagues have effectively employed intermediate causal models to teach students about electrical circuits.

Intermediate constructs such as facets also serve to link other related concepts using the underlying pathophysiology of disease, which results in the development of more coherent problem explanations. Intermediate constructs can then be viewed as "gluing" constructs that serve to link data (e.g., patient findings) to hypothesis (e.g., diagnosis) into whole units.

The conceptual structures that experts generate are typically coherent, because most pertinent patient data are accounted for with a single explanation, for instance, a diagnosis. Furthermore, when the recall of clinical information by physicians is analyzed for coherence, it is found that the greater the coherence, the better the quality of the diagnosis. In contrast, the recalls and explanations given by intermediate subjects (e.g., medical students) are only partly coherent. Although intermediate subjects may possess the relevant knowledge, their explanations are often fragmented (i.e., focused on isolated pieces of information, with few links between clinical components). However, for prior knowledge to be used efficiently, it needs to be organized in a way that generates specific inferences that link patient data to the hypothesis (e.g., diagnosis, treatment) in a direct manner.

These findings have implications for the development of information systems. People at all levels of expertise will likely be users of information systems. However, given the varieties of expertise and skill levels in medicine, it is important that information systems be

responsive to different developmental characteristics of the users. Consider the subexpert presented in Figure 3. This physician does not generate intermediate constructs effectively, because already at the second piece of information, the physician suggests the hypothesis of emphysema. That is, the subexpert begins to explore a path that is too specific at this point. A biomedical information system tuned to the expertise level of the physician-user may use reminders at the facet-level hypotheses (e.g., pericardial disease) that are also related to the patient data and showing him or her the specific links that tie the data to the construct (e.g., exposing the underlying pathophysiology). For instance intermediate constructs can be represented as a network of concepts where each concept is linked to other relevant concepts. In this way, the subexpert would be presented in an explicit manner by the system what the expert does implicitly.

The Nature of Conceptual Understanding

Claim 8: Flaws in conceptual knowledge are composed of systematic patterns of misunderstanding with complex etiologies, which lead to specific types of errors in prediction and explanation.

Information technologies are transforming the practice of medicine and having a substantial impact on medical education. Some educators have suggested that the introduction of new technologies greatly diminishes the need for medical students to acquire basic-science knowledge, since that knowledge will be readily accessible via computers.[104] This argument obscures the fact that conceptual understanding is predicated on prior knowledge. Computers are a potentially powerful instructional medium, but mere access to information is not likely to promote meaningful learning. In addition, one needs to have a certain domain mastery to interpret information effectively. It is likely that the best clinical judgment will continue to require a broad understanding of basic-science knowledge.[105]

As discussed previously, conceptual understanding necessitates the acquisition of generative knowledge. The generativity of knowledge is indicated by the breadth of applicability across problem types, including situations that were never previously encountered. For example, rote learning from a medical textbook is not sufficient to produce conceptual change (i.e., change in conceptual understanding). How can medical education foster such understanding?

We are interested in understanding how biomedical knowledge can support reasoning and justify action in diverse cognitive tasks in medicine and health care domains—tasks undertaken by trained professionals,

professionals-in-training and lay people. Specifically, we are interested in characterizing dimensions of conceptual competence that entail "acquiring and retaining a network of concepts and principles about some domain that accurately represent key phenomena and their interrelationships and that can be engaged flexibly when pertinent to accomplish diverse, sometimes novel objectives."[106] Several challenges are involved in fostering conceptual competence, including developing a coherent understanding of a set of related concepts constituting a subset of a particular domain, learning to apply these problems, and the transfer of this knowledge to a range of situations.[107] Understanding a set of concepts, as reflected in the ability of a person to generate an explanation, does not guarantee that he or she will be able to apply this knowledge in a particular-decision-making context.[108] Students who learn a concept in a particular context (e.g., a particular clinical case) have difficulty applying it in other contexts.

Numerous studies in mathematical and science domains have documented, among students, widespread misconceptions that persist even after instruction.[109] Feltovich et al.[106,110] have carried out a series of studies pertaining to cardiovascular physiology based on an elaborate framework for complex conceptual analysis. A clear example of their work is reflected in a misconception related to congestive heart failure.[110] In this syndrome the heart's effectiveness as a pump can diminish greatly, and as a result the rate of blood flow slows dramatically. The misconception that was expressed by more than 60 percent of first- and second-year medical students, and by some medical practitioners, suggested that heart failure is caused by the heart getting too big, which in turn stretches the cardiac muscle fibers. The force of contraction is determined by mechanical/anatomic factors and activational (energetic) factors. The primary cause of congestive heart failure is activational, whereas the misconception emphasizes the mechanical overstretching as the cause of heart failure. The authors identified several component misconceptions that tend to interact and create robust misunderstandings of important phenomena.

Employing a similar framework, we evaluated the ability of first-year medical students to use the concept of ventilation/perfusion matching in the lungs to explain a problem of a patient with an embolus obstructing blood flow.[82] The results showed systematic misconceptions by students in developing a pathophysiologic model of the problem. The subjects demonstrated an inability to coordinate events in the right and left lungs and in the dysfunctional and func-

tional regions of lung tissue. Students frequently were not able to conceptualize the cardiopulmonary system as a closed system, with an event in one region propagating effects throughout other regions of the lung. We also studied how students, physicians, and physiologists understood cardiac output and its regulation. The study revealed that subjects at more advanced levels of expertise demonstrated a more differentiated understanding of concepts pertaining to cardiovascular physiology, as evidenced by their explanations and predictions.[111] However, we also observed faulty reasoning in some of our experts.

The study documented four significant misconceptions. The most pervasive misconception was a confounding of venous compliance and resistance. Compliance refers to the expandability of a vessel and its capacity to store blood. Venous resistance is primarily a function of the radius of the vessel. An increase in compliance increases the volume storage capacity of the vessel and therefore decreases venous return. Likewise, an increase in resistance impedes the flow of blood and slows venous return.

The misconception was evident in the responses of six subjects, including two fourth-year students and two residents. These subjects indicated that venous resistance would diminish compliance and therefore increase blood flow. The sources of erroneous knowledge and beliefs that contributed to this misconception include the belief that venous resistance and compliance are completely interdependent; the notion that the large veins are storage vessels (medical students are frequently taught that veins are "storers of blood") when they are, in fact, resistance vessels; and a malprioritization of causal factors resulting in a misinterpretation of the primary effect of resistance.

The investigation of flaws in understanding and misconceptions has shown that errors are often the result of deeply held beliefs about fundamental biomedical processes. These beliefs can have detrimental effects on performance (e.g., causing errors) in a variety of tasks, ranging from diagnosis to treatment and management. Typically, errors of perceptual or memory nature can be identified more easily than errors of faulty understanding, since the latter can remain hidden from observed performance. Their effects can also be more damaging. Because computerized systems, such as electronic medical records, impose a structure on a user's problem solving, they can conceal such misconceptions and flaws.

What are the consequences of misconceptions for practicing physicians? Evidence suggests that health care professionals (e.g., physicians and nurses) can correctly diagnose problems or make appropriate decisions without a full understanding of the reasons for their actions.[4,112] However, the errors they make may be simple procedural slips or they may be based on more serious faulty conceptual knowledge; either of these can contribute substantially to unacceptable medical practices. For example, an inadequate understanding of the determinants of venous return may lead to inappropriate treatment strategies for left heart failure. Furthermore, even technical errors (e.g., missing steps in procedures) may be grounded in conceptual misunderstandings of the type discussed here.

Problem Solving and Reasoning

Claim 9: Expertise is associated with differential use of reasoning strategies in problem solving.

Data-driven reasoning, in which data triggers the solution to a problem in a relatively effortless, automated way, constitutes the primary reasoning strategy used by experts in routine clinical problems. In contrast, hypothesis-driven reasoning (from hypothesis to data) is used in complex and uncertain problems.

Problem solving has been a major area of research in medical cognition since the 1970s, when Elstein et al.,[28] motivated by the work of Newell and Simon,[27] published a seminal series of studies on how physicians solved clinical problems. Research in many different problem-solving tasks extended and deepened our understanding of the cognitive processes involved in medical problem solving. A major focus of the research that followed focused on the strategies that physicians and medical trainees used to solve clinical diagnostic problems.[38,113,114] It was found that the use of different reasoning strategies during problem solving was associated with differences in expertise.

In studies of problem solving, two main patterns of reasoning strategies were identified[62,64]—data-driven and hypothesis-driven reasoning (see Figure 4). The patterns of data-driven and hypothesis-driven reasoning are used differentially by novices and experts. Experts tend to use data-driven reasoning, which depends on the physician's possessing a highly organized knowledge base about the patient's disease (including sets of signs and symptoms). Because of their lack of substantive knowledge or their inability to distinguish relevant from irrelevant knowledge, novices and intermediates use more hypothesis-driven reasoning, resulting often in very complex reasoning patterns. The fact that experts and novices reason differently suggests that they might reach different conclusions (e.g., decisions or understandings) when solving medical problems. Similar patterns of

Figure 4 Diagrammatic representation of data-driven *(top)* and hypothesis-driven reasoning *(bottom)*. From the presence of vitiligo, a prior history of progressive thyroid disease, and examination of the thyroid *(clinical findings on the left)*, the physician reasons forward to conclude the diagnosis of myxedema *(on the right)*. However, the anomalous finding of respiratory failure, which is inconsistent with the main diagnosis, is accounted for as a result of a hypometabolic state of the patient, in a backward-directed fashion. COND indicates a conditional relation; CAU, a causal relation; RSLT, a resultive relation.
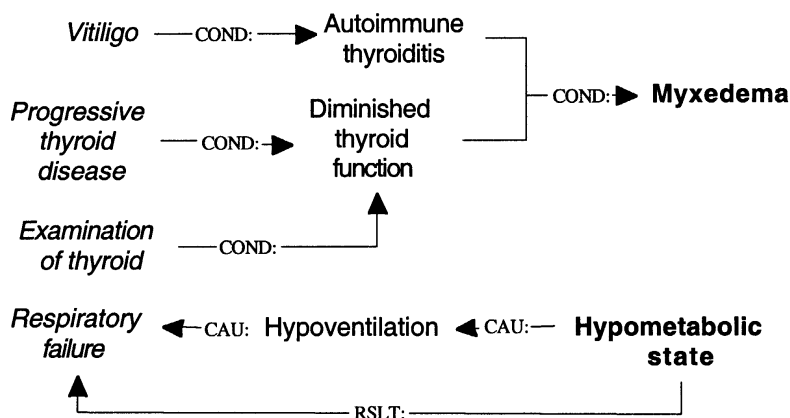
reasoning have been found in other domains.[88] Because of their extensive knowledge base and the high-level inferences they make, experts typically skip steps in their reasoning.

Although experts typically use data-driven reasoning during clinical performance, this type of reasoning sometimes breaks down, and the expert has to resort to hypothesis-driven reasoning. In the absence of adequate domain knowledge, data-driven reasoning is highly prone to error, since built-in checks on the legitimacy of inferences are lacking. Pure data-driven reasoning is successful only in constrained situations, in which knowledge of a problem can result in a chain of inferences from the initial problem statement to the problem solution, as illustrated in Figure 4.

In contrast, hypothesis-driven reasoning is slower and may make heavy demands on working memory, because of the need to keep track of such things as goals and hypotheses. Hypothesis-driven reasoning is, therefore, most likely to be used when domain knowledge is inadequate or the problem is complex. Hypothesis-driven reasoning is a type of "weak" method, i.e., a strategy that is used when the problem solver either lacks knowledge about the domain or is highly uncertain about the solution. However, it is typically used with data-driven reasoning rather than alone.[64]

Studies have shown that the pattern of data-driven reasoning breaks down in conditions of case complexity, unfamiliarity with the problem, and uncertainty.[114] These conditions include the presence of "loose ends" in explanations, when a particular piece of information remains unaccounted for and isolated from the overall explanation. Loose ends trigger explanatory processes that work by a person's hypothesizing a disease, for instance, and trying to fit

the loose ends into the hypothesis, in hypothesis-driven reasoning fashion. The presence of loose ends may foster learning, as the person searches for an explanation for them. For instance, a medical student or physician may encounter a sign or a symptom in a patient and look for information that accounts for the finding by searching for similar cases seen in the past, reading a specialized medical book, or consulting a domain expert.

In some circumstances, however, the use of data-driven reasoning may lead to a heavy cognitive load. This refers to the existence of substantial information that competes for few cognitive resources, creating a burden on working memory.[115,116] For instance, giving students problems to solve when they are being trained in the use of problem-solving strategies places a heavy load on their cognitive resources and diminishes their ability to focus on the task. The reason is that the students have to divide their cognitive resources (e.g., attention and memory) between learning the problem-solving method and solving the problems.

It has been found that when subjects used a strategy based on data-driven reasoning, they were more able to acquire a schema for a problem. In addition, other characteristics associated with expert performance have been observed, such as making fewer moves to reach a solution. However, when subjects used hypothesis-driven reasoning strategies, their problem-solving performance suffered.[115,116]

Understanding the patterns of diagnostic reasoning used by clinicians at varying levels of expertise can inform the development of decision support technology and the use of such technology in the ebb and flow of clinical work. To this end, we have been pursuing the study of how technology affects such reasoning patterns.[49]

Conclusions

Cognitive science can provide important insights into the nature of the cognitive processes involved in human–computer interaction and thereby improve the application of medical information systems by addressing the knowledge, memory, and strategies used in a variety of cognitive activities. Medical cognition is a discipline that draws on theories and methods from cognitive science. Until recently, the connection between medical informatics and medical cognition has been somewhat remote. In this paper, we have surveyed literature on aspects of medical cognition and provide a set of claims that we consider to be important in medical informatics.

Although great technologic achievements can be made in the development of medical systems, the design of such systems should be informed by cognitive constraints that users impose on the systems with which they interact. Knowledge of the way people process information is crucial to the successful use of information systems. It is often acknowledged that the health informatics community needs to develop a common, uniform, and comprehensive approach to the representation of medical information to help in its utilization. This would require collaboration to arrive at a consensus regarding what knowledge should be presented and how. Cognitive research shows that the use of this information by physicians changes as a function of their background, developmental level, and expertise. Standardization, although necessary, has its limitations. The question remains as to the extent to which particular systems can be flexibly adapted for specific persons, situations, and conditions. Cognitive science has much to offer in this regard.

A number of areas of research that are relevant to medical cognition have not been included in this primer. Some of these have been presented in other sources.[117,118] They cover a vast range of issues, such as the role of perception in clinical diagnosis, decision making in natural settings, education and learning and decision-support systems, and the nature and development of collaboration. A great deal of research[41,119] on how experts and novices interpret visual material (e.g., x-rays and skin lesions) in making diagnoses has been addressed only briefly.

The study of decision making in medicine, for instance, has a long tradition of research,[11] which has recently turned to the study of decision making in complex naturalistic settings.[82,120] A second field pertains to education and learning in real-world environments,[111] including assessment,[121] which are likely to have an important role in informatics-based education of health care professionals.[122] This field also involves the study of collaborative decision making. With the increasing costs of health care, the need for collaboration among physicians and institutions has become more pressing.[123] Fortunately, great advances in collaborative tools—ranging from simple ones like e-mail to more complex ones like telemedical systems—have made collaboration more realizable. Research into these issues has already begun to produce promising results.[6] Ongoing work in these fields of inquiry should provide more insight into the cognitive problems that arise in medical informatics.

References ■

1. Greenes RA, Shortliffe EH. Medical informatics: an emerging academic discipline and institutional priority. JAMA. 1990; 263:1114–20.
2. Patel V, Kaufman D. Medical informatics and the science of cognition. J Am Med Inform Assoc. 1998;5:493–502.
3. Hunt E. Will We Be Smart Enough? A Cognitive Analysis of the Coming Workforce. New York: Russell Sage, 1995.
4. Kaufman DR, Kushniruk AW, Yale JF, Patel VL. Conceptual knowledge and decision strategies in relation to hypercholesterolemia and coronary heart disease. Int J Med Inform. 1999; 55:159–77.
5. Kushniruk AW, Kaufman RD, Patel VL, Lévesque Y, Lottin P. Assessment of computerized patient record system: a cognitive approach to evaluation of an emerging medical technology. MD Comput. 1996;13:406–15.
6. Patel VL, Kaufman DR, Allen VG, Shortliffe EH, Cimino JJ, Greenes RA. Toward a framework for computer-mediated collaborative design. Methods Inf Med. 1999;38(3):158–76.
7. Shackel B. Usability: context, framework, definition, design, and evaluation. In: Shackel B, Richardson S (eds). Human Factors for Informatics Usability. Cambridge, UK: Cambridge University Press, 1991.
8. Blois M. Clinical judgment and computers. N Engl J Med. 1980;303:192–7.
9. Patel VL, Cytryn KN, Shortliffe EH, Safran C. The collaborative health care team: the role of individual and group expertise. Teaching Learning Med. 2000;12:117–32.
10. Chapman GB, Sonnenberg FS. Decision Making in Health Care: Theory, Psychology, and Applications. Cambridge, UK: Cambridge University Press, 2000.
11. Dowie J, Elstein A. Professional Judgment: A Reader in Clinical Decision Making. Cambridge, UK: Cambridge University Press, 1988.
12. Hmelo CE. Computer assisted instruction in health professions education: a review of the published literature. J Educ

Technol Syst. 1989/90;18:83–101.

13. Lillehaug SI, Lajoie SP. AI in medical education: another grand challenge for medical informatics. Artif Intell Med. 1998;12:197–225.

14. Mishler EG. The Discourse of Medicine: The Dialectics of Medical Interviews. Norwood, NJ: Ablex, 1984.

15. Atkinson P. Medical Talk and Medical Work: The Liturgy of the Clinic. London, UK: Sage, 1995.

16. Sarangi S, Roberts C. Talk, work and institutional order. discourse in medical mediation and management settings. In: Heller M, Watts RJ (eds). Language, Power and Social Process. Berlin, Germany: Mouton de Gruyter, 1999.

17. Kohn LT, Corrigan JM, Donaldson MS (eds). To Err is Human: Building a Safer Health System. Washington, DC: National Academy Press, 1999.

18. Bogner MS. Human Error in Medicine. Hillsdale, NJ: Lawrence Erlbaum, 1994.

19. Riese W. The structure of Galen's diagnostic reasoning. Bull NY Acad Med. 1968;44:778–91.

20. Osler W. Aequanimitas: With Other Addresses to Medical Students, Nurses and Practitioners of Medicine. 2nd ed., with 3 additional addresses. Philadelphia, Pa.: P. Blakiston's Son, 1906.

21. Leaper DJ, Horrocks JC, Staniland JR, de Dombal FT. Computer-assisted diagnosis of abdominal pain using "estimates" provided by clinicians. BMJ. 1972;2:350–4.

22. Ericsson KA, Simon HA. Protocol analysis: verbal reports as data (rev ed). Cambridge, Mass.: MIT Press, 1993.

23. Ledley RS, Lusted LB. Reasoning foundations of medical diagnosis. Science. 1959;130:9–21.

24. Feinstein AR. An analysis of diagnostic reasoning: the domain and disorders of clinical macrobiology. Yale J Biol Med. 1973;46:264–83.

25. Rimoldi HJA. The test of diagnostic skills. J Med Educ. 1961; 36:73–9.

26. Kleinmuntz B. Formal Representation of Human Judgment. New York: Wiley, 1968.

27. Newell A, Simon HA. Human Problem Solving. Englewood Cliffs, NJ: Prentice Hall, 1972.

28. Elstein AS, Shulman LS, Sprafka SA. Medical problem solving: an analysis of clinical reasoning. Cambridge, Mass.: Harvard University Press, 1978.

29. Shortliffe EH. The adolescence of AI in medicine: Will the field come of age in the 90s? Artif Intell Med. 1993;5:93–106.

30. Gorry GA. Modeling the diagnostic process. J Med Educ. 1970;45:293–302.

31. Pauker SG, Gorry GA, Kassirer JP, Schwartz WB. Towards the simulation of clinical cognition: taking a present illness by computer. Am J Med. 1976;60:981–96.

32. Miller RA, Pople HE, Myers DJ. Internist-I, an experimental computer-based diagnostic for general internal medicine. In: Clancey WJ, Shortliffe EH (eds). Readings in Medical Artificial Intelligence. Reading, Mass.: Addison-Wesley, 1984:190–209.

33. Buchanan BG, Shortliffe EH. Rule-based expert systems: the MYCIN experiments of the Stanford Heuristic Programming Project. Reading, Mass.: Addison-Wesley, 1984.

34. Patil RS, Szolovits P, Schwartz WB. Causal understanding of patient illness in medical diagnosis. In: Clancey WJ, Shortliffe EH (eds). Readings in Medical Artificial Intelligence. Reading, Mass.: Addison-Wesley, 1984:339–60.

35. Clancey WJ. Acquiring, representing and evaluating a competence model of diagnostic strategy. In: Chi MTH, Glaser R, Farr MJ (eds). The Nature of Expertise. Hillsdale, NJ: Lawrence Erlbaum, 1988:343–418.

36. Clancey WJ, Letsinger R. NEOMYCIN: reconfiguring a rule-based expert system for application to teaching. In: Clancey WJ, Shortliffe EH (eds). Readings in Medical Artificial Intelligence: The First Decade. Reading, Mass.: Addison-Wesley; 1984:361–81.

37. Feltovich PJ, Johnson PE, Moller JH, Swanson DB. LCS: The role and development of medical knowledge in diagnostic expertise. In: Clancey WJ, Shortliffe EH (eds). Readings in Medical Artificial Intelligence: The First Decade. Reading, Mass.: Addison-Wesley; 1984:275–319.

38. Patel VL, Groen GJ. Knowledge-based solution strategies in medical reasoning. Cognitive Sci. 1986;10:91–116.

39. Norman G, Brooks LR, Rosenthal D, Allen SW, Muzzin LJ. The development of expertise in dermatology. Arch Dermatol. 1989;125:1063–8.

40. Chi MTH, Feltovich PJ, Glaser R. Categorization and representation of physics problems by experts and novices. Cognitive Sci. 1981;5:121–52.

41. Lesgold A, Rubinson H, Feltovich P, Glaser R, Klopfer D, Wang Y. Expertise in a complex skill: diagnosing x-ray pictures. In: Chi MTH, Glaser R, Farr MJ (eds). The Nature of Expertise. Hillsdale, NJ: Lawrence Erlbaum, 1988:311–42.

42. Patel VL, Arocha JF. Medical expertise and cognitive aging. In: Park D, Morrell RW, Shifren K (eds). Processing of Medical Information in Aging Patients: Cognitive and Human Factors Perspectives. Mahwah, NJ: Lawrence Erlbaum, 1999.

43. Anderson JR, Reder LM, Simon HA. Applications and misapplications of cognitive psychology to mathematics education. Pittsburgh, Pa.: Carnegie-Mellon University, 1995.

44. Bechtel W, Abrahamsen A, Graham G. The life of cognitive science. In: Bechtel W, Graham G, Balota DA (eds). A Companion to Cognitive Science. Malden, Mass.: Blackwell, 1998:2–104.

45. Durso F, Nickerson RS, Schvaneveldt RW, Dumais ST, Lindsay DS, Chi MTH. Handbook of Applied Cognition. New York: Wiley, 1999:881.

46. Hunt E. Cognitive science: definition, status, and questions. Ann Rev Psychol. 1989;40:603–29.

47. Ericsson A, Smith J. Prospects and limits of the empirical study of expertise: an introduction. In: Ericsson A, Smith J (eds). Toward a General Theory of Expertise: Prospects and Limits. New York: Cambridge University Press, 1991:1–38.

48. Miller GA. The magical number seven, plus or minus two: some limits on our capacity for processing information. Psychol Rev. 1956;63:81–97.

49. Patel VL, Kushniruk AW, Yang S, Yale J-F. Impact of computer-based patient record system on data collection, knowledge organization, and reasoning. J Am Med Inform Assoc. 2000;7:569–85.

50. Patel VL, Arocha JF, Kushniruk AW. EMR: re-engineering the organization of health information. Methods Inf Med, in press.

51. Bransford JD, Brown AL, Cocking RR. How People Learn: Brain, Mind, Experience, and School. Expanded edition. Committee on Developments in the Science of Learning. Washington, DC: Commission on Behavioral and Social Sciences and Education/National Academy Press, 2000.

52. Bainbridge L. Multiplexed VDT display systems. In: Weir GRS, Alty JL (eds). Human–Computer Interaction and Complex Systems. London, UK: Academic Press, 1991:189–210.

53. Ericsson KA, Krampe RT, Tesch-Romer C. The role of deliberate practice in the acquisition of expert performance. Psychol Rev. 1993;100:363–406.

54. Ericsson KA, Kintsch W. Long-term working memory.

Psychol Rev. 1995;102:211–45.

55. Charness N. Expertise in chess: the balance between knowledge and search. In: Ericsson KA, Smith J (eds). Toward a General Theory of Expertise: Prospects and Limits. New York: Cambridge University Press, 1991:39–63.

56. Falasconi S, Stefanelli M. A library of medical ontologies. In: Mars NJI (ed). Proceedings of ECAI-94 (Workshop on Comparison of Implemented Ontologies); Aug 8–12, 1994; Amsterdam, The Netherlands. 1994:81–91.

57. Patel VL, Ramoni M. Cognitive models of directional inference in expert medical reasoning. In: Feltovich P, Ford K, Hoffman R (eds). Expertise in Context. Menlo Park, Calif.: AAAI Press/MIT Press, 1997:67–99.

58. Evans DA, Gadd CS. Managing coherence and context in medical problem-solving discourse. In: Evans DA, Patel VL (eds). Cognitive Science in Medicine: Biomedical Modeling. Cambridge, Mass.: MIT Press, 1989:211–55.

59. Chandrasekeran B, Smith JW, Sticklen J. Deep models and their relation to diagnosis. Artif Intell Med. 1989;1:29–40.

60. Pople H. Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnosis. In: Szolovits P (ed). Artificial Intelligence in Medicine. Boulder, Colo.: Westview Press, 1982:119–90.

61. Patel VL, Arocha JF. Cognitive models of clinical reasoning and conceptual representation. Methods Inf Med. 1995;34:1–10.

62. Patel VL, Arocha JF, Kaufman DR. Diagnostic reasoning and expertise. Psychol Learning Motivation. 1994;31:137–252.

63. Patel VL, Evans DA, Kaufman DR. Cognitive framework for doctor-patient interaction. In: Evans DA, Patel VL (eds). Cognitive Science in Medicine: Biomedical Modeling. Cambridge, Mass.: MIT Press, 1989:253–308.

64. Patel VL, Groen GJ. The general and specific nature of medical expertise: a critical look. In: Ericsson A, Smith J (eds). Toward a General Theory of Expertise: Prospects and Limits. New York: Cambridge University Press, 1991:93–125.

65. Kushniruk AW, Patel VL, Cimino JJ, Barrows RA. Cognitive evaluation of the user interface and vocabulary of an outpatient information system. Proc AMIA Annu Fall Symp. 1996:22–6.

66. Gorry GA. Strategies for computer-aided diagnosis. Math Biosci. 1968;2:293–318.

67. Josephson J, Josephson S. Abductive inference, computation, philosophy, technology. Cambridge, UK: Cambridge University Press, 1994.

68. Chandrasekaran B. Generic tasks in knowledge-based reasoning: high-level building blocks for expert system design. IEEE Expert. 1986;1:23–30.

69. Gott SP, Lesgold AM. Competence in the workplace: how cognitive performance models and situated instruction can accelerate skill acquisition. In: Glaser R (ed). Advances in Instructional Psychology: Educational Design and Cognitive Science. Mahwah, NJ: Lawrence Erlbaum, 2000.

70. Cole M, Engeström Y. A cultural-historical approach to distributed cognition. In: Salomon G (ed). Distributed Cognitions: Psychological and Educational Considerations. Cambridge, Mass.: Cambridge University Press; 1993:1–46.

71. Woods DD. Observations from studying cognitive systems in context. Proceedings of the 16th Annual Conference of the Cognitive Science Study. Hillsdale, NJ: Lawrence Erlbaum, 1994:961–4.

72. Pea RD. Practices of distributed intelligence and designs for education. In: Salomon G (ed). Distributed Cognitions: Psychological and Educational Considerations. Cambridge, Mass.: Cambridge University Press, 1993:47–87.

73. Hutchins E. Cognition in the Wild. Cambridge, Mass.: MIT Press, 1996.

74. Lave J. Teaching as learning, in practice. Mind Culture Activity. 1996;3:149–64.

75. Salomon G, Perkins D. Individual and social aspects of learning. In: Pearson PD, Iran-Nejad A (eds). Review of Research in Education. Washington, DC: American Educational Research Association, 1998:1–24.

76. Greeno JG, Collins AM, Resnick LB. Cognition and learning. In: Berliner DC, Calfee RC (eds). Handbook of Educational Psychology. New York: MacMillan, 1996:15–46.

77. Patel VL (ed). Distributed and collaborative cognition in medicine: implications for system development [special issue]. Artif Intell Med. 1998;12(2):93–195.

78. Timpka T, Westergren V, Hallberg N, Forsum U. Study of situation-dependent clinical cognition: a meta-analysis and preliminary method. Methods Inf Med. 1997;36:44–50.

79. Suchman LA. Plans and situated actions: the problem of human machine communication. In: Pea R, Brown JS (eds). Learning in doing: social, cognitive, and computational perspectives. Cambridge, UK: Cambridge University Press, 1987.

80. Vicente KJ. Cognitive work analysis: toward safe, productive, and healthy computer-based work. Mahwah, NJ: Lawrence Erlbaum, 1999:392.

81. Hutchins E. The social organization of distributed cognition. In: Resnick LB, Levine JM, Teasley SD (eds). Perspectives on Socially Shared Cognition. Washington, DC: American Psychological Association, 1991:283–307.

82. Patel VL, Kaufman DR, Magder SA. The acquisition of medical expertise in complex dynamic environments. In: Ericsson KA (ed). The Road to Excellence: The Acquisition of Expert Performance in the Arts and Sciences, Sports and Games. Hillsdale, NJ: Lawrence Erlbaum, 1996:127–65.

83. Norman DA. Cognitive artifacts. In: Carroll JM (ed). Designing Interaction: Psychology at the Human–Computer Interface. Cambridge, Mass.: Cambridge University Press, 1991:17–38.

84. Zhang J. Representations of health concepts: a cognitive perspective. In: Chute C (ed). Proceedings of IMIA Working Group 6 Conference on Medical Concept Representation and Natural Language Processing; Dec 16–19, 1999; Phoenix, Ariz.; pp. 44-49.

85. Polson PG, Lewis C, Rieman J, Wharton C. Cognitive walkthroughs: a method for theory-based evaluation of user interfaces. Int J Man–Machine Studies. 1992;36:741–73.

86. Ericsson YA, Smith J. Toward a general theory of expertise: prospects and limits. New York: Cambridge University Press, 1991.

87. deGroot AD. Thought and Choice in Chess. The Hague, The Netherlands: Mouton, 1965.

88. Larkin JH, McDermott J, Simon HA, Simon DP. Expert and novice performances in solving physics problems. Science. 1980;208:1335–42.

89. Sloboda J. Musical expertise. In: Ericsson KA, Smith J (eds). Toward a general theory of expertise: prospects and limits. New York: Cambridge University Press, 1991:153–71.

90. Patel VL, Groen GJ. Developmental accounts of the transition from student to physician: some problems and suggestions. Med Educ. 1991;25:527–35.

91. Patel VL, Groen GJ, Scott HS. Biomedical knowledge in explanations of clinical problems by medical students. Med Educ. 1988;22:398–406.

92. Patel VL, Groen GJ, Patel YC. Cognitive aspects of clinical performance during patient workup: the role of medical expertise. Adv Health Sci Educ. 1997;2:95–114.

93. Arocha JF, Patel VL, Patel YC. Hypothesis generation and the coordination of theory and evidence in novice diagnostic rea-

soning. Med Decis Making. 1993;13:198–211.

94. Schmidt HG, Boshuizen HP. On the origin of intermediate effects in clinical case recall. Memory Cogn. 1993;21:338–51.

95. Johnson PE, Duran AS, Hassebrock F, et al. Expertise and error in diagnostic reasoning. Cognitive Sci. 1981;5:235–83.

96. Saariluoma P. Apperception and restructuring in chess players' problem solving. In: Gilhooly K, Keane MTG, Logie RH, Erdos G (eds). Lines of Thinking. London, UK: Wiley, 1991:41–71.

97. Robertson RJ, Glines LA. The phantom plateau returns. Percept Motor Skills. 1985;61:55–64.

98. Strauss S, Stavy R. U-shaped Behavior Growth. New York: Academic Press, 1982.

99. Clark HH. Using Language. Cambridge, UK: Cambridge University Press, 1996.

100. Kintsch W. Comprehension: A Paradigm for Cognition. New York: Cambridge University Press, 1998:461.

101. Joseph G-M, Patel VL. Domain knowledge and hypothesis generation in diagnostic reasoning. Med Decis Making. 1990; 10:31–46.

102. Coiera E. Intermediate depth representations. Artif Intell Med. 1992;4:431–45.

103. White BY. ThinkerTools: Causal models, conceptual change, and science education. Cognition and Instruction. 1993;10:1–100.

104. Cavazos LF. Basic science studies: their purpose in medical education. J Med Educ. 1984;59:763.

105. Prokop DJ. Basic science and clinical practice: How much will a physician need to know? In: Marston RQ, Jones RM (eds). Medical Education in Transition. Princeton, NJ: The Robert Wood Johnson Foundation, 1992:51–7.

106. Feltovich PJ, Spiro RJ, Coulson RL. Learning, teaching, and testing for complex conceptual understanding. In: Frederiksen N, Mislevy RJ, Bejar II (eds). Test Theory for a New Generation of Tests. Hillsdale, NJ: Lawrence Erlbaum, 1993:181–217.

107. The Cognition and Technology Group at Vanderbilt. Adventures in anchored instruction: lessons from beyond the Ivory Tower. In: Glaser R (ed). Advances in Instructional Psychology: Educational Design and Cognitive Science. Mahwah, NJ: Lawrence Erlbaum, 2000:35–99.

108. Kaufman DR, Patel VL. Progressions of mental models in understanding circulatory physiology. In: Singh I, Parasuraman R (eds). Human Cognition: A Multidisciplinary Perspective. New Delhi, India: Sage India, 1998.

109. Eylon B, Linn MC. Research perspectives in science education. Rev Educ Res. 1988;58:251–301.

110. Feltovich PJ, Spiro R, Coulson RL. The nature of conceptual understanding in biomedicine: the deep structure of complex ideas and the development of misconceptions. In: Evans DA, Patel VL (eds). Cognitive Science in Medicine: Biomedical Modeling. Cambridge, Mass.: The MIT Press, 1989:113–72.

111. Patel VL, Kaufman DR, Arocha JF. Conceptual change in the biomedical and health sciences domain. In: Glaser R (ed). Advances in Instructional Psychology: Educational Design and Cognitive Science. Mahwah, NJ: Lawrence Erlbaum, 2000:329–92.

112. Leprohon J, Patel VL. Decision making strategies for telephone triage in emergency medical services. Med Decis Making. 1995;15:240–53.

113. Patel VL, Groen GJ, Frederiksen CH. Differences between students and physicians in memory for clinical cases. Med Educ. 1986;20:3–9.

114. Patel VL, Groen GJ, Arocha JF. Medical expertise as a function of task difficulty. Memory and Cognition. 1990;18:394–406.

115. Chandler P, Sweller J. Cognitive load theory and the format of instruction. Cognition and Instruction. 1991;8:293–32.

116. Sweller J, Chandler J. Evidence for cognitive load theory. Cognition and Instruction. 1991;8:351–62.

117. Clancey WJ, Shortliffe EH. Readings in Medical Artificial Intelligence: The First Decade. Reading, Mass.: Addison-Wesley, 1984.

118. Evans DA, Patel VL. Cognitive science in medicine: biomedical modeling. Hillsdale, NJ: Lawrence Erlbaum, 1989.

119. Norman GR, Coblentz CL, Brooks LR, Babcock CJ. Expertise in visual diagnosis: a review of the literature. Acad Med. 1992;67:S78–S83.

120. Klein A, Orasanu J, Calderwood R, Zsambok CE. Decision making in action: models and methods. Norwood, NJ: Ablex, 1993138–47.

121. Glaser R. Toward new models for assessment. Int J Educ Res. 1990;14:475–83.

122. Glaser R, Ferguson EL, Vosniadou S. Cognition and the design of environments for learning: approaches in this book. In: Vosniadou S, DeCorte E, Glaser R, Mandl H (eds). International Perspectives on the Design of Technology-supported Learning Environments. Mahwah, NJ: Lawrence Erlbaum, 1996.

123. Shortliffe EH, Patel VL, Cimino JJ, Barnett GO, Greenes RA. A study of collaboration among medical informatics research laboratories. Artif Intell Med. 1998;12:97–123.