Thinking and Reasoning in Medicine

Vimla L. Patel, Columbia University, NY patel@dbmi.columbia.edu

Jose F. Arocha, University of Waterloo, Canada jfarocha@healthy.uwaterloo.ca

Jiajie Zhang, University of Texas at Houston, TX Jiajie.Zhang@uth.tmc.edu

Citation:

Patel, V. L., Arocha, J. F., & Zhang, J. (2004, in press). Thinking and reasoning in medicine. In: Keith Holyoak: *Cambridge Handbook of Thinking and Reasoning.* Cambridge, UK: Cambridge University Press.

Key Words

Medical reasoning, medical knowledge, diagnosis, abductive reasoning, case-based reasoning, knowledge organization, forward reasoning, hypothesis testing, backward reasoning, expertise, intermediate effect, technology and thinking, distributed reasoning, external representation, representational effect, representation determinism, electronic medical records, human error, medical error, actions, biases, heuristics.

What is Medical Reasoning?

Medical reasoning describes a form of qualitative inquiry that examines the cognitive (thought) processes involved in making medical decisions. Clinical reasoning, medical problem solving, diagnostic reasoning, and decision-making are all terms used in a growing body of literature that examines how clinicians make clinical decisions. Medical cognition refers to studies of cognitive processes, such as perception, comprehension, decision making, and problem solving in medical practice itself or in tasks representative of medical practice. These studies use subjects who work in medicine, including medical students, physicians, and biomedical scientists. The study of medical reasoning has been the focus of much research in cognitive science and artificial intelligence in medicine. Medical reasoning involves an inferential process for making diagnostic or therapeutic decisions or understanding the pathology of a disease process. On the one hand, medical reasoning is basic to all higher-level cognitive processes in medicine such as problem solving and medical text comprehension. On the other hand, the structure of medical reasoning is itself the subject of considerable scrutiny. For example, the directionality of reasoning in medicine has been an issue of considerable controversy in medical cognition, medical education and artificial intelligence in medicine. Conventionally, we can partition medical reasoning into clinical and biomedical or basic science reasoning. These are some of the central themes that constitute this chapter.

Early Research on Medical Problem Solving and Reasoning

Medical cognition is a subfield of cognitive science devoted to the study of cognitive processes in medical tasks. Studies of medical cognition include analyses of performance in “real-world” clinical tasks as well as in experimental tasks. 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 (Osler, 1906).

Human information processing research has typically focused on the individual. The dual focus on in-depth task analysis and on the study of human performance is a central feature of a cognitive science approach.

There have been 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 (Dowie and Elstein, 1988). It compares the performance of a physician with the mathematical model by focusing on reasoning "fallacies" and biases inherent in human clinical decision-making (Leaper et.al 1972). In contrast, the information-processing approach focuses on the description of cognitive processes in reasoning tasks and the development of cognitive models of performance, typically relying on protocol analysis (Ericsson and Simon, 1993) and other observational techniques.

Systematic investigations of medical expertise began more than 40 years ago with the research by Ledley and Lusted (1959) on clinical inquiries. They proposed a two-stage model of clinical reasoning involving a hypothesis-generation stage followed by a hypothesis-evaluation stage, where the latter stage was amenable to formal decision analytic techniques. Probably the earliest empirical studies of medical reasoning can be traced to the work of Rimoldi (1961) 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 to 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 (1972) on problem solving, research in information-processing psychology accelerated dramatically. Problem solving was conceived of as search in a problem space in which a problem solver was viewed as selecting an option (e.g., a hypothesis or an inference) or performing an operation (from a set of possible operations) in moving toward a solution or a goal state (e.g., diagnosis or treatment plan). (See Novick & Bassok, Chap. 11, for a discussion of problem solving.) This conceptualization had an enormous impact in both cognitive psychology and artificial intelligence research. It also led to rapid advances in medical reasoning and problem solving research, as exemplified by the seminal work of Elstein, Shulman, & Sprafka (1978). 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 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 (AI), particularly, expert systems technology. AI in medicine and medical cognition mutually influenced each other in a number of ways, which included 1) providing a basis for developing formal models of competence in problem-solving tasks; 2) elucidating the structure of medical knowledge and providing important epistemological distinctions, and 3) characterizing productive and less-productive lines of reasoning in diagnostic and therapeutic tasks. Gorry (1973) conducted a series of studies comparing a computational model of medical problem solving 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 and colleagues (1976) capitalized on some of the insights of Gorry's earlier work and developed the Present Illness Program (PIP), a program designed to take the history of a patient with edema. Several of the 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 partially based on studies of clinical problem solving.

Medical expert consultation systems such as Internist (Miller, Pople & Myers, 1982) and MYCIN (Shortliffe, 1976) introduced the 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. MYCIN also highlighted the difference between medical problem solving and the cognitive dimensions of medical explanation. Clancey's work (1984,1985) in GUIDON and NEOMYCIN was particularly influential in the evolution of models of medical cognition. Clancey endeavored to reconfigure MYCIN in order 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 upon.

Feltovich and colleagues (Feltovitch, Johnson, Moller et al.1984), drawing on models of knowledge representation from medical artificial intelligence, characterized fine-grained differences in knowledge organization between subjects with 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 (1986), incorporating distinctions introduced by Clancey, studied the knowledge-based solution strategies of expert cardiologists as evidenced by their pathophysiological 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 toward a complete diagnosis (i.e., reasoning from data to hypothesis). In contrast, subjects who misdiagnosed or partially diagnosed the patient problem used a backward reasoning strategy. These research findings presented a challenge to the hypothetico-deductive model of reasoning as espoused by Elstein et.al (1978), which did not differentiate expert from non-expert reasoning strategies.

Much of the early research in the study of reasoning in domains such as medicine was carried out in laboratory or experimental settings. In more recent times, a shift then occurred toward examining cognitive issues in naturalistic medical settings, such as medical teams in intensive care units (Patel, Kaufman, & Magder, 1996), anesthesiologists working in surgery (Gaba, 1992), nurses providing emergency telephone triage (Leprohon & Patel, 1995), and reasoning with technology by patients (Patel, Kuyshniruk etal, 2002) in the health care system. This research was been informed by work in the area of dynamic decision-making (Salas & Klein, 2000*)*, complex problem solving (Frensch & Funke, 1995), human factors (Hoffman & Deffenbacher, 1992; Vicente & Rasmussen, 1990), and cognitive engineering (Rasmussen, Pejtersen, & Goodstein, 1994). Such studies, conducted in the workplace, reshaped our views of human thinking by shifting the onus of cognition from being the unique province of the individual to being distributed across social and technological contexts.

Models of Medical Reasoning

The traditional view of medical reasoning has been to treat diagnosis as similar to the scientist's task of making a discovery or engaging in scientific experimentation (see Dunbar & Fugelsang, Chap. 29). Coherent with this view of science is the assumption that diagnostic inference follows a hypothetico-deductive process of reaching conclusions by testing hypothesis based on clinical evidence. Within a cognitive perspective, as we saw previously, this view of the diagnostic process in medicine was first proposed in the influential work of Elstein, Shulman, and Sprafka (1978). Such view of medical reasoning as hypothetico-deductive has been challenged from various points, empirical research and philosophical discourse, as we will see in this section.

Toward a Model of Reasoning in Medicine: Induction, Deduction, and Abduction

It is generally agreed upon that there are two basic forms of reasoning. One is deductive reasoning (see Evans, Chap. 6), which consists of deriving a particular valid conclusion from a set of general premises, and the other is inductive reasoning (see Sloman & Lagnado, Chap. 3), which consists of deriving a likely general conclusion from a set of particular statements. However, reasoning in the “real world” does not appear to fit neatly into any of these basic types. For this reason, a third form of reasoning has been recognized, where deduction and induction are inter-mixed. This was termed “abductive reasoning” by Pierce (1955).

Basically, all theories of medical reasoning characterize diagnosis as an abductive, cyclical, process of generating possible explanations (i.e., identification of a set of hypotheses that are able to account for the clinical case on the basis of the available data) and testing those explanations (i.e., evaluation of each generated hypothesis on the basis of its expected consequences) for the abnormal state of the patient at hand (Elstein, et. al, 1978; Kassirer, 1989; Joseph & Patel, 1990; Ramoni et al., 1992). Traditional accounts of medical reasoning have described diagnostic process in a way that is independent of the underlying structure of the domain knowledge. These accounts simply make the assumption that some domain of knowledge exists and that all of the hypotheses needed to explain a problem are available when the diagnostic process starts.

Within this generic framework, various models of diagnostic reasoning may be constructed. Following Patel and Ramoni (1997), we could distinguish between two major models of diagnostic reasoning: *heuristic classification* (Clancey, 1985) and *cover and differentiate* (Eshelman, 1988). However, these models can be seen as special cases of a more general model: the *select and test* model, where the processes of hypothesis generation and testing can be characterized in terms of four types of inferences (Peirce, 1955): abstraction, abduction, deduction, and induction. The first two inference types drive hypothesis generation while latter two types drive hypothesis testing. During *abstraction*, data are filtered according to their relevance for the problem solution and chunked in schemas representing an abstract description of the problem at hand (e.g., abstracting that an adult male with hemoglobin concentration less than 14d/gl is an anemic patient). Following this, hypotheses that could account for the current situation are related through a process of *abduction*, characterized by a "backward flow" of inferences across a chain of directed relations which identify those initial conditions from which the current abstract representation of the problem originates. This provides tentative solutions to the problem at hand by way of hypotheses. For example, knowing that disease *A* will cause symptom *b*, abduction will try to identify the explanation for *b*, while deduction will forecast that a patient affected by disease *A* will manifest symptom *b*: both inferences are using the same relation along two different directions. These three types of reasoning in medicine is described in a paper by Patel and Ramoni (1997).

In the testing phase, hypotheses are incrementally tested according to their ability to account for the whole problem, where *deduction* serves to build up the possible world described by the consequences of each hypothesis. This kind of reasoning is customarily regarded as a common way of evaluating diagnostic hypotheses (Kassirer, 1989; Patel, Evans, & Kaufman, 1989; Joseph & Patel, 1990; Patel, Arocha, & Kaufman, 1994). As predictions are derived from hypotheses, they are matched to the case through a process of *induction*, where a prediction generated from a hypothesis can be matched with one specific aspect of the patient problem. The major feature of induction is, therefore, the ability to rule out those hypotheses whose expected consequences turn out to be not in agreement with the patient problem. This is because there is no way to logically confirm a hypothesis, but we can only disconfirm or refute it in the presence of contrary evidence. This evaluation process closes the testing phase of the diagnostic cycle. Moreover, it determines which information is needed in order to discriminate among hypotheses and hence which information has to be collected.

Hypothesis Testing and Clinical Reasoning

Although a model such as one presented above can be used to account for large part of the medical diagnostic process, empirical literature has pointed to various strategies of diagnostic reasoning that underscore the relative importance of deduction, induction, or abduction. In their seminal work, Elstein and colleagues (Elstein et al., 1978) studied the problem solving processes of physicians by drawing on then contemporary methods and theories of cognition. Their view of problem solving had a substantial influence on both studies of medical reasoning and medical education. They were the first to use experimental methods and theories of cognitive science to investigate clinical competency. Their research findings 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). First, physicians generated a small set of hypotheses very early in the case, as soon as the first pieces of data became available. Second, they were selective in the data they collected, focusing only on the relevant data. Third, physicians made use of the hypothetico-deductive process, which consisted of four stages: cue acquisition, hypothesis generation, cue interpretation, and hypothesis evaluation. Cues in the clinical case led to the generation of a few selected hypotheses, where each cue was interpreted as positive, negative or non-contributory to each hypothesis generated. Then each hypothesis was evaluated for consistency with the cues. Using this framework, these investigators were unable to find differences between superior physicians (as judged by their peers) and other physicians (Elstein et al., 1978).

Forward-driven and Backward-driven Reasoning

Later, Patel and Groen (1986) studied knowledge-based solution strategies of expert cardiologists as evidenced by their pathophysiological explanations of a complex clinical problem. The results indicated that subjects who accurately diagnosed the problem employed a forward-oriented (data-driven) reasoning strategy—using patient data to lead toward a complete diagnosis (i.e., reasoning from data to hypothesis). This was in contrast to subjects who misdiagnosed or partially diagnosed the patient problem, who tended to use a backward or hypothesis-driven reasoning strategy. The results of this study presented a challenge to the hypothetico-deductive model of reasoning as espoused by Elstein and colleagues (1978) which did not differentiate expert from non-expert reasoning strategies.

A hypothesis for reconciling these seemingly contradictory results is that forward reasoning is used in clinical problems in which the physician has ample experience. However, when reasoning through unfamiliar or difficult cases, physicians resort to backward reasoning since their knowledge base does not support a pattern-matching process. To support this explanation, Patel, Groen, and Arocha (1990) looked for the conditions under which forward reasoning breaks down. Cardiologists and endocrinologists were asked to solve diagnostic problems both in cardiology and in endocrinology. They showed that under conditions of case complexity and uncertainty, the pattern of forward reasoning was disrupted. More specifically, the breakdown occurred when non-salient cues in the case were tested for consistency against the main hypothesis, even in subjects who had generated the correct diagnosis. Otherwise, the results supported previous studies in that subjects with accurate diagnoses used pure forward reasoning.

If forward reasoning breaks down when case complexity is introduced, then experts and novices should reason differently because routine cases for experts would not be so for less-than-expert subjects. Investigating clinical reasoning in a range of contexts of varying complexity (Patel & Groen, 1991; Patel, Arocha Kaufman, 1994), the authors found that novices and experts have different patterns of data-driven and hypothesis-driven reasoning. As before, experts used data-driven reasoning, which depends on the physician possessing a highly organized knowledge base about the patient's disease (including sets of signs and symptoms). Furthermore, due to their extensive knowledge base and the high level inferences they make, experts typically skip steps in their reasoning. In contrast, because of their lack of substantive knowledge or their inability to distinguish relevant from irrelevant knowledge, less than expert subjects (novices and intermediates) used more hypothesis-driven reasoning, resulting often in very complex reasoning patterns. Similar patterns of reasoning have been found in other domains (Larkin et al., 1980).

The fact that experts and novices reason differently suggests that they might reach different conclusions (e.g., decisions or understandings) when solving medical problems. Although data-driven reasoning is highly efficient, it is often error prone in the absence of adequate domain knowledge, since there are no built-in checks on the legitimacy of the inferences that a person makes. Pure data-driven reasoning is only successful in constrained situations, where one's knowledge of a problem can result in a complete chain of inferences from the initial problem statement to the problem solution. In contrast, hypothesis-driven reasoning is slower and requires high memory load, because one has to keep track of such things as goals and hypotheses. It is therefore most likely to be used when domain knowledge is inadequate or the problem is complex. Hypothesis-driven reasoning is an exemplar of a *weak method* of problem solving in the sense that it is used in the absence of relevant prior knowledge and when there is uncertainty about problem solution. In problem-solving terms, strong methods engage knowledge whereas weak methods refer to general strategies. Weak does not necessarily imply ineffectual in this context.

Studies also showed that data-driven reasoning can break down due to uncertainty (Patel, Groen, & Arocha, 1990). These conditions include the presence of "loose ends" in explanations, where some particular piece of information remains unaccounted for and isolated from the overall explanation. Loose ends trigger explanatory processes that work by hypothesizing a disease, for instance, and trying to fit the loose ends within it, in a 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 a physician may encounter a sign or a symptom in a patient problem and look for information that may account for the finding, by searching for similar cases seen in the past, reading a specialized medical book, or consulting a domain expert. (See Chi & Ohlsson, Chap. 12, for a discussion of such complex forms of learning.)

However, in some circumstances, the use of data-driven reasoning may lead to a heavy cognitive load. For instance, when students are given problems to solve while they are being trained in the use of problem solving strategies, the situation produces a heavy load on cognitive resources which may diminish students' ability to focus on the task. The reason is that students have to share cognitive resources (e.g., attention, memory) between learning the problem-solving method and learning the content of the material. Research (Sweller, 1988) suggests that when subjects use a strategy based on the use of data-driven reasoning, they are more able to acquire a schema for the problem. In addition, other characteristics associated with expert performance were observed, such as a reduced number of moves to the solution. However, when subjects used a hypothesis-driven reasoning strategy, their problem solving performance suffered. The study of medical reasoning has been summarized in a series of articles (e.g. Patel et al., 1994; Patel et al., 2002) and papers in edited volumes (Clancey & Shortliffe,1984; Szolovits.1982).

The Role of Similarity in Diagnostic Reasoning

The fact that physicians make use of forward reasoning in routine cases suggests a type of processing that is fast enough to be able to lead to the recognition of a set of signs and symptoms in a patient and generate a diagnosis based on such recognition. Most often this has been interpreted as a type of specific-to-general reasoning (e.g., reasoning from an individual case to a clinical schema or prototype). However, consistent with the model of abductive reasoning, some philosophers (Schaffner, 1986) and empirical researchers (Norman & Brooks, 1997) have supported an alternative hypothesis, which consists of specific-to-specific reasoning. That is, experts also use knowledge of specific instances (e.g., particular patients with specific disease presentations) to interpret particular cases, rather than relying only on general clinical knowledge (Kassirer & Kopelman, 1990).

Brooks and colleagues (Brooks, Norman, & Allen, 1991; Norman and Brooks, 1997) have argued that clinicians make use of specific instances in order to compare and interpret a current clinical case. In such studies, mostly involving visual diagnosis—such as X-rays, dermatological slides, and electrocardiograms—it has been shown that specific similarity to previous cases accounts for about 30%

of diagnoses made (see Goldstone & Son, Chap. 1; Rips & Medin, Chap. 2). Furthermore, errors made by experts in identifying abnormalities in images are affected by the prior history of the patient. That is, if the prior history of the patient mentioned a possible abnormality, expert physicians more often identified abnormalities in the images even when none were there, which also supports the effect of specific past cases on the interpretation of a current case.

In pursuing their explanation, Norman and colleagues (Norman and Brooks, 1997) argued against the hypothesis that expert physicians diagnose clinical cases by “analyzing” signs and symptoms and developing correspondences between those signs, symptoms and diagnoses, as traditional cognitive research in medical reasoning suggests. They suggest instead the “non-analytic” basis for medical diagnosis, where diagnostic reasoning is characterized by the unanalyzed retrieval of a similar case previously seen in medical practice to interpret the current case: a kind of exemplar-based or case-based reasoning. They. This discussion has its counterpart in the psychology of categorization, where two accounts have been proposed: either categorization works by a reliance on prototypes or by exemplars (Rips & Medin, Chap. 2).

Exemplar-based thinking is certainly a fundamental aspect of human cognition. There is ample evidence of the conditions where reasoning by analogy to previous cases is used (Gentner & Holyoak, 1997; Holyoak & Thagard, 1997). Furthermore, given the complexity of natural reasoning in a highly dense knowledge domain such as medicine, it is highly likely that more than one type of reasoning is actually employed. Seen in this light, the search for a single manner in which clinicians diagnose clinical problems may not be a reasonable goal. The inherent adaptability of humans to different kinds of knowledge domains, situations, problems, and cases may call for the use of a variety of reasoning strategies, which is what, after all, the notion of abductive medical reasoning has tried to formalize (Patel & Ramoni, 1997. Alongside with rule-based and prototype reasoning, a model of clinical reasoning may allow for case-based, non-analytical reasoning, where similarity between particulars may be the main cognitive mechanism. A reason for the variety of strategies used in actual diagnostic problems may be found in the inherent organization of medical knowledge.

Reasoning and the Nature of Medical Knowledge

Although a motivation for looking at medical reasoning was to establish its relationship with reasoning in other fields, such as science, the prevalent view in the philosophy of medicine (Blois, 1988) has been that medical knowledge has an extremely complex organization, requiring the use of different reasoning strategies than those used in other more formal scientific disciplines, such as physics. Disciplines such as physics, chemistry, and some subfields of biology, are said to be *horizontally* organized, where these domains are characterized by the construction of causal relations among concepts and by the application of general principles to specific instances Blois (1988). By this, it is meant that such scientific fields are organized in a hypothetico-deductive manner where particular statements are generated from general statements, and where causality plays a major role. This type of reasoning, in which one connects one concept to another by forming causal networks, has been called “horizontal” reasoning (Blois, 1988). These philosophers have argued that causal reasoning does not play such an important role in the medical domain. They argue, instead, that reasoning in medicine requires “vertical“ thinking. This kind of reasoning makes more use of the analogy than the reasoning typically found in other scientific domains. In this view, the medical disciplines, notably clinical medicine, are organized *vertically*, and reasoning by analogy (see Holyoak, Chapter 4) plays a more important role than causal reasoning. Based on such a distinction, it has been further been argued that reasoning in the physical sciences and reasoning in the biomedical sciences are of different kind.

In particular, it has been argued that reasoning in physical sciences can be, to some extent, conceptualized as a "deductive systematization of a broad class of generalizations under a small number of axioms", but this characterization cannot be applied to the biomedical sciences. The latter are characterized by what Shaffner (1986, p. 68) calls "a series of overlapping interleaved temporal models", which are based on familiarization with shared exemplars to a much greater degree than is necessary in the physical sciences. Shaffner’s characterization, unlike that of Blois, applies to both biomedical research and clinical medicine. In biomedical research, an organism such as a *Drosophila*, for instance, is used as an exemplar embodying a given disease mechanism, which by analogy applies to other organisms, including humans. In the clinical sciences, the patient is seen as an exemplar to which generalizations based on multiple overlapping models are applied from diseases and the population of similar patients.

In the empirical research on medical reasoning the distinction between reasoning from cases versus reasoning from prototypes has not been established. Medical knowledge consists of two categories of knowledge: clinical knowledge, including knowledge of disease processes and associated findings; and basic science knowledge, incorporating subject matter such as biochemistry, anatomy, and physiology. Basic science or biomedical knowledge is supposed to provide a scientific foundation for clinical reasoning. The conventional view is that basic science knowledge can be seamlessly integrated into clinical knowledge analogous to the way that learning the rules of the road can contribute to one’s mastery of driving a car. In this capacity, a particular piece of biomedical knowledge could be automatically elicited in a range of clinical contexts and tasks in more or less the same fashion.

Knowledge Organization and Changes in Directionality

Following Blois (1988) and Schaffner (1986), it can be argued that the way medical knowledge is organized can be a determinant factor explaining why experts do not use the hypothetico-deductive method of reasoning. Maybe the medical domain is too messy to allow its neat partitioning and deductive use of reasoning strategies. Although the theory of reasoning in medicine is basically a theory of expert knowledge, reaching the level of efficient reasoning of the expert clinician reflects the extended continuum of training and levels of reasoning performance (Thibodeau, Hardiman, Dufresne, & Mestre, 1989; Chi, Bassok, Lewis, Glaser, & Reiman, 1989). This continuum also points to the particular nature of medical knowledge and its acquisition.

Changes have been described in this process that serve to characterize the various phases medical trainees go through to become expert clinicians. An important characteristic of this process is the *intermediate effect*. This refers to the fact that, although it seems reasonable to assume that performance improves with training or time on task, there appear to be particular transitions, in which subjects exhibit a certain drop in performance. This is an example of what is referred to as non-monotonicity in the developmental literature (Strauss & Stavy, 1982) and is also observed in skill acquisition. The symptom is a learning curve or developmental pattern that is shaped like either a U or an inverted U, as illustrated in Figure 1. In medical expertise development, intermediates' performance reflects the degradation in reasoning that results from the acquisition of knowledge through a time during which such knowledge is not well-organized and irrelevant associations abound in the intermediate's knowledge-base. In contrast, the novice's knowledge-base is too sparse containing very few associations whereas the expert's knowledge-base is well pruned of the irrelevancies that characterize intermediates. It should be noted that not all intermediate performance is non-monotonic; for example, on some global criteria such as diagnostic accuracy, there appears to be a steady improvement.

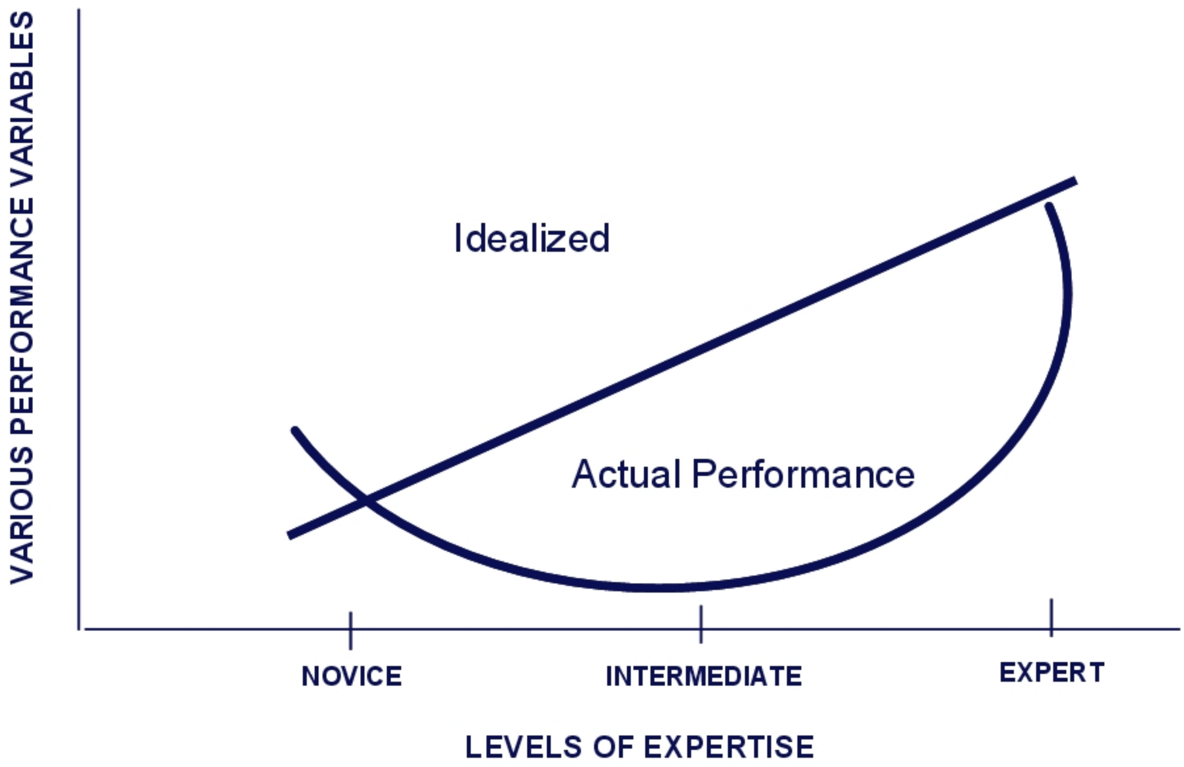

Figure 1: Idealized representation of the "Intermediate effect". The straight line gives a commonly assumed representation of performance development by level of expertise. The curved, U-shaped, line represents the actual development from novice to expert. The Y-axis may represent performance variables, such as the number of errors made, irrelevant concepts recalled, conceptual elaborations, or number of extraneous hypotheses generated in a variety of tasks.

The intermediate effect occurs with many tasks and at various levels of expertise. The tasks vary from comprehension of clinical cases and explanation of clinical problems to problem solving to generating laboratory data. The phenomenon may be due to the fact that intermediates have acquired an extensive body of knowledge, but have not yet reorganized this knowledge in a functional manner. Thus intermediate knowledge has a sort of network structure that results in considerable search, which makes it more difficult for intermediates to set up structures for rapid encoding and selective retrieval of information (Patel& Groen, 1991). In contrast, expert knowledge is finely tuned to perform various tasks and experts can readily filter out irrelevant information using their hierarchically organized schemata. The difference is reflected both in the structural organization of knowledge and the extent to which it is proceduralized to perform different tasks.

Schmidt and Boshuizen (1993) reported that intermediate non-monotonicity recall effects disappear by using short exposure times (about 30 seconds), which suggests that under time-restricted conditions, intermediates cannot engage in extraneous search. While a novice's knowledge base is likely to be sparse and an expert's knowledge base is intricately interconnected, the knowledge base of an intermediate possesses many of the pieces of knowledge in place, but lacks the extensive connectedness of an expert. Until this knowledge becomes further consolidated, the intermediate is more likely to engage in unnecessary search. Whether this knowledge, painfully acquired during medical training, is really necessary for clinical reasoning has been a focus of intensive research and great debate. If expert clinicians do not explicitly use underlying biomedical knowledge, does that mean that it is not necessary? Or could it be simply the case that this knowledge remains "dormant" until is really needed? This raises an important question of whether expert medical knowledge is "deep" or "shallow".

Causal Reasoning in Medicine

The differential role of basic science knowledge (e.g., physiology and biochemistry) in solving problems of varying complexity and the differences between subjects at different levels of expertise (Patel et. al, 1994) have been a source of controversy in the study of medical cognition (Patel & Kaufman, 1995) as well as in medical education and artificial intelligence. As expertise develops, the disease knowledge of a clinician becomes more dependent on clinical experience and clinical problem solving is increasingly guided by the use of exemplars and analogy, and becomes less dependent on a functional understanding of the system in question. However, an in-depth conceptual understanding of basic science plays a central role in reasoning about complex problems and is also important in generating explanations and justifications for decisions.

AI researchers were confronted with similar problems in extending the utility of systems beyond their immediate knowledge base. Biomedical knowledge can serve different functional roles depending on the goals of the system. Most models of diagnostic reasoning in medicine can be characterized as being "shallow." For instance, a "shallow" medical expert system (e.g., MYCIN and INTERNIST) reasons by relating observations to intermediate hypotheses that partition the problem space, and further associating intermediate hypotheses with diagnostic hypotheses. This is consistent with the way physicians appear to reason. There are however, other medical reasoning system models that propose a "deep" mode of reasoning as a main mechanism. Chandrasakeran et al. (1989) characterize a deep system as one, which embodies a causal mental model of bodily function and malfunction, similar to the models used in qualitative physics (Bobrow, 1985). Systems such as MDX-2 (Chandrasakeran et al., 1989) or QSIM (Kuipers, 1987) have explicit representations of structural components and their relations, the functions of these components (in essence their purpose), and their relationship to behavioral states.

To become licensed physicians, medical trainees undergo a lengthy training process that includes the learning of biomedical sciences, including biochemistry, physiology, anatomy, and others. It has been pointed out the apparent contradiction between this type of training and the absence of “deep” biomedical knowledge during expert medical reasoning. In order to account for such apparent inconsistency, Boshuizen and Schmidt (1992) proposed a learning mechanism, *knowledge encapsulation.* Knowledge encapsulation is a learning process, which involves the subsumption of biomedical propositions and their interrelations in associative clusters, under a small number of higher-level clinical propositions with the same explanatory power. Through exposure to clinical training, biomedical knowledge presumably becomes integrated with clinical knowledge. Biomedical knowledge can be “unpacked” when needed, but is not used as a first line of explanation.

Boshuizen and Schmidt (1972) cite a wide range of clinical reasoning and recall studies that support this kind of learning process. Of particular importance is the well-documented finding that with increasing levels of expertise, physicians produce explanations at higher levels of generality, using fewer and fewer biomedical concepts while producing consistently accurate responses. The intermediate effect can also be accounted for as a stage in the encapsulation process in which a trainee's network of knowledge has not yet become sufficiently differentiated, thus resulting in more extensive processing of information.

Knowledge encapsulation provides an appealing account of a range of developmental phenomena in the course of acquiring medical expertise. However, the integration of basic science in clinical knowledge is a rather complex process, and encapsulation is likely to be only part of the knowledge development process. Basic science knowledge plays a different role in different clinical domains. For example, clinical expertise in perceptual domains, such as dermatology and radiology, necessitates a relatively robust model of anatomical structures that is the primary source of knowledge for diagnostic classification. In other domains, such as cardiology and endocrinology, basic science knowledge has a more distant relationship with clinical knowledge. The misconceptions evident in physicians' biomedical explanations would argue against their having well developed encapsulated knowledge structures in which basic science knowledge could easily be retrieved and applied when necessary.

The results of research into medical problem solving are consistent with the idea that clinical medicine and biomedical sciences constitute two distinct and not completely compatible "worlds", with distinct modes of reasoning and quite different ways of structuring knowledge (see Patel, Arocha, & Kaufman, 1994). Clinical knowledge is based on a complex taxonomy that relates disease symptoms to underlying pathology. In contrast, biomedical sciences are based on general principles defining chains of causal mechanisms. Thus, learning to explain how a set of symptoms is consistent with a diagnosis may be very different from learning how to explain what causes a disease. (See Buehner & Cheng, Chap. 5, for a discussion of causal reasoning.)

The notion of the progression of mental models (White & Frederiksen, 1990) has been used as an alternative framework for characterizing the development of conceptual understanding in biomedical contexts. Mental models are dynamic knowledge structures that are composed to make sense of experience and to reason across spatial and/or temporal dimensions. An individual's mental models provide predictive and explanatory capabilities of the function of a given system. The authors employed the progression of mental models to explain the process of understanding increasingly sophisticated electrical circuits. This notion can be used to account for differences between novices and experts in understanding circulatory physiology, describing misconceptions (Kaufman & Patel, 1994) and explaining the generation of spontaneous analogies in causal reasoning (Kaufman, Patel, & Magder, 1996).

Running a mental model is a potentially powerful form of reasoning but it is also cognitively demanding. It may require an extended chain of reasoning and the use of complex representations. It is apparent that skilled individuals learn to circumvent long chains of reasoning and chunk or compile knowledge across intermediate states of inference (Chandrasekaran, 1994; Newell, 1990). This results in shorter, more direct, inferences that are stored in long-term memory and are directly available to be retrieved in the appropriate contexts. Chandrasekaran (1994) refers to this sort of knowledge as *compiled causal knowledge*. This term refers to knowledge of causal expectations that people compile directly from experience and partly by chunking results from previous problem-solving endeavors (Kaufman & Patel, 1994). The goals of the individual and the demands of recurring situations largely determine which pieces of knowledge get stored and used. When a physician is confronted with a similar situation, she can employ this compiled knowledge in an efficient and effective manner. The development of compiled knowledge is an integral part of the acquisition of expertise.

The idea of compiling declarative knowledge bears a certain resemblance to the idea of knowledge encapsulation. However, the claim differs in two important senses. The process of compiling knowledge is not one of subsumption or abstraction, and the original knowledge (uncompiled mental model) may no longer be available in a similar form (Kuipers & Kassirer, 1984). The second difference is that mental models are composed dynamically out of constituent pieces of knowledge rather than pre-stored unitary structures. The use of mental models is somewhat opportunistic and the learning process is less predictable. The compilation process can work in reverse as well. That is to say, discrete cause-and-effect relationships can be integrated into a mental model as a student reasons about complex physiological processes.

Errors and Medical Reasoning

According to the report from the Institute of Medicine (Kohn, Corrigan, & Donaldson, 1999), medical error is the eighth leading cause of death in the US, ahead of deaths due to motor vehicle accidents, breast cancer, or AIDS. Cognitive mechanisms, such as mistakes of reasoning and decision making and action slips of skilled performance, are the major factors contributing to medical errors. A cognitive taxonomy is essential for the understanding, explanation, and prediction of medical errors and for the development of interventions to reduce medical errors. Based on the definition and the preliminary taxonomy by Reason (1992) and the action theory by Norman (1986), Zhang, Patel, Johnson, & Shortliffe (2004, in review) developed a cognitive taxonomy for human errors in medicine.

A Cognitive Taxonomy of Medical Errors

One critical step towards understanding the cognitive mechanisms of various errors in medical reasoning is to categorize the errors along cognitively meaningful dimensions. Reason (1992) defines human error as a failure of achieving the intended outcome in a planned sequence of mental or physical activities. He divides human errors into two major categories: (1) slips that result from the incorrect execution of a correct action sequence and (2) mistakes that result from the correct execution of an incorrect action sequence. Norman's theory of action (Norman, 1986) decomposes a human activity into seven stages.

Based on Reason's definition of human error and Norman's action theory, Zhang and colleagues developed a cognitive taxonomy. Under this taxonomy errors are divided into slips and mistakes, just like Reason's two main categories. Then slips are divided into execution slips and evaluation slips. Execution slips include goal, intention, action specification, and action execution slips, whereas evaluation slips include perception, interpretation and evaluation slips. Similarly, mistakes can also be divided into execution mistakes that include goal, intention, action specification, and action execution mistakes, and evaluation mistakes that include perception, interpretation and evaluation mistakes. This taxonomy can cover major types of medical errors, because a medical error is a human error in an action and any action goes through the seven stages of the action cycle. Most reasoning and decision-making errors in medicine are under the category of mistakes in the taxonomy. They are due to incorrect or incomplete knowledge

Reasoning and Decision Making Mistakes in Medicine

In the cognitive taxonomy, goal and intention mistakes are mistakes about declarative knowledge, which is knowledge about factual statements and propositions, such as "Motrin is a pain reliever and fever reducer". Action specification mistakes and action execution mistakes are mistakes about procedural knowledge, which is knowledge about procedures and rules, such as “give 1 tsp Motrin to a child per dosage up to 4 times a day if the child has fever or toothache and the weight of the child is 24-35 lbs.”

Goal mistakes and intention mistakes are caused by many complex factors such as incorrect knowledge, incomplete knowledge, and misuse of knowledge, biases, faulty heuristics, and information overload. For example, neglect of base rate information could result in incorrect diagnosis of a disease. This is a well-documented finding in human decision making (Tversky & Kahneman, 1974; Kahneman & Frederick, Chap. 10). As another example, the goal of “treating the disease as pneumonia” could be a mistake if it is a misdiagnosis based on incomplete knowledge (e.g., without x-ray images). Intention mistakes can be caused by similar factors, such as the following example: A physician treating a patient with oxygen set the flow control knob between 1 and 2 liters per minute, not realizing that the scale numbers represented discrete, rather than continuous, settings. As a result, the patient did not receive any oxygen. This is a mistake due to incomplete knowledge. The use of heuristics is another common source of goal and intention mistakes. A heuristic that is often used is the reliance on disease schemata during clinical diagnosis. Disease schemata are knowledge structures that have been formed from previous experience with diagnosing diseases and contain information about relevant and irrelevant signs and symptoms. When physicians and medical students diagnose patients, they tend to rely on their schemata and base their reasoning on the apparent similarity of patient information with these schemata, instead of a more objective analysis of patient data. The schemata that are used in diagnosis often guide future reasoning about the patient, affecting what tests are run and how data are interpreted. Arocha and Patel (1995) found that medical students and trainees maintained their initial hypotheses, even if subsequent data were contradictory. Therefore, if the initial hypothesis is wrong, errors in diagnosis and treatment are likely to occur. Preliminary presentation of the patient (e.g., signs and symptoms), then, becomes very important, because it can suggest strongly held hypotheses (i.e., lead to the use of schemata).

Action specification and action execution mistakes are procedural mistakes that can be caused by many factors such as lack of correct rules, over-generalized application of good rules, misapplication of good rules, encoding deficiencies in rules, and the dissociation between knowledge and rules. For example, over-generalized application of good rules can cause an error because the condition part of a condition-action rule could be misidentified and mismatched, thus causing the firing of the action part of the rule. Procedural mistakes caused by encoding efficiencies of action rules are usually due to the evolving nature of the rules and unforeseeable conditions that cannot be encoded in the rules. A good rule may be misused because the user may have incorrect or incomplete knowledge about the condition of the rule in a specific context. The knowledge of a rule and the knowledge of how to use a rule are not always automatically linked together without extensive practice. This dissociation, due to the lack of experience and practiced skills, may also lead to action execution mistakes.

Perception mistakes can be caused by expectation-driven processing. What we perceive is a function of the input and our expectations. This mechanism is what allows us to read sloppy handwriting, or recognize degraded images. However, our expectations can also lead to misperceptions. Interpretation mistakes are the incorrect interpretation of feedback caused by incorrect or incomplete knowledge. For instance, suppose that an intravenous infusion pump, which is a device often used in critical care environments to give medications, indicates readiness to begin infusing medications using a steady green light and indicates the infusion is in progress by flashing the green light. If the device user does not know the meaning of the steady green light, he or she may incorrectly interpret it as an indication that the infusion has begun. Another source of interpretation mistake is generations of different interpretations and treatment procedures from the same evidence. An action evaluation mistake occurs when incorrect knowledge or incomplete knowledge leads a person to erroneously judge the completion or incompletion of a goal.

Medical Reasoning and Decision Research

Decision-making is central to medical activity. Although health-care professionals are generally highly proficient decision-makers, their erroneous decisions have become the source of considerable public scrutiny (Kohn et al., 1999).

Decisions involve the application of reasoning to select some course of action that achieves the desired goal (see LeBoeuf & Shafir, Chap. 9). Hastie (2001) has identified three components of decision making: (a) choice options and courses of actions; (b) beliefs about objective states, processes and events in the world, including outcomes states and means to achieve them; and (c) desires, values or utilities that describe the consequences associated with the outcomes of each action-event combination. In this process, reasoning plays a major role. In this context, a major thrust of research has been the study of hypothesis testing, which has been widely studied in the medical domain. Such research has shown the pervasiveness of confirmation bias, which is evidenced by the generation of a hypothesis and the subsequent search for evidence consistent with the hypothesis, often leading to the failure to adequately consider alternative diagnostic possibilities. This bias may result in a less than thorough investigation with possible adverse consequences for the patient. A desire to confirm one's preferred hypothesis may moreover contribute to increased inefficiency and costs by ordering additional laboratory tests that will do little to revise one's opinion, providing largely redundant data (Chapman & Elstein, 2000).

Health care team decision-making is the rule rather than the exception in medicine. Naturalistic decision making (NDM) is concerned with the study of cognition in "real-world" work environments that are often dynamic (e.g., rapidly changing) (Klein et al., 1993; Lipshitz et al., 2001). The majority of this research combines conventional protocol analytic methods with innovative methods designed to investigate reasoning and behavior in realistic settings (Woods, 1993; Rasmussen et al., 1994). The study of decision making in the work context necessitates an extended cognitive science framework beyond typical characterizations of knowledge structures, processes, and skills to include modulating variables such as stress, time pressure, and fatigue as well as communication patterns in team performance.

Among the issues investigated in NDM are understanding how decisions are jointly negotiated and updated by participants differing substantially in their areas of expertise (e.g., pharmacology, respiratory medicine); how the complex communication process in these settings occurs; what role technology plays in mediating decisions and how it affects reasoning; and what the sources of error are in the decision making process.

Research by Patel, Kaufman, and Magder (1996) studied decision-making in a medical intensive care unit (ICU) with the particular objective of describing jointly negotiated decisions, communication processes and the development of expertise. Intensive care decision-making is characterized by a rapid serial evaluation of options leading to immediate action, where reasoning is schema-driven in a forward direction towards action with minimal inference or justification. However, when patients do not respond in a manner consistent with the original hypothesis, then the original decision comes under scrutiny. This strategy can result in a brainstorming session in which the team retrospectively evaluates and reconsiders the decision and considers possible alternatives. In such circumstances, various patterns of reasoning are used to evaluate alternatives in these 'brainstorming' sessions. These include probabilistic reasoning, diagnostic reasoning, and biomedical causal reasoning. Supporting decision-making in clinical settings necessitates an understanding of communication patterns.

In summary, although traditional approaches to decision making looked at decisions as choosing among known alternatives, real-world decision making is best investigated by a naturalistic approach in which reasoning is constrained by dynamic factors, such as stress, time pressure, risk, and team interactions. Looking at medical reasoning in social and collaborative settings is even more important when information technologies are part of the ebb and flow of clinical work.

Reasoning and Medical Education

The failures and successes of reasoning strategies and skills can be traced back to their sources: education. There is evidence suggesting that the way physicians reason follows from the way they have been educated. Medical education in North America as well as in the rest of the world has followed a similar path: from practice-based training to an increasingly scientific training.

Motivated by the increasing importance of basic scientific knowledge in the context of clinical practice, problem-based learning (PBL) was developed on the premise that not only should physicians possess the ordered and systematic knowledge of science, but also they should *think* like scientists during their practices. Consistent with this idea, an attempt was made to teach hypothetico-deductive reasoning to medical students, as an attempt to provide an adequate structure to medical *problem solving*. After all, this was the way scientists were supposed to make discoveries.

Based on cognitive research in other knowledge-domains, some researchers argued, however, that the hypothetico-deductive method might not be the most efficient way of solving clinical problems. To investigate how the kind of training medical students receive affected their reasoning patterns, Patel, Groen, and Norman (1993) looked at the problem-solving processes of students in two medical schools with different modes of instruction, classical and problem-based. They found that students in the problem-based curriculum reasoned in a way that was consistent with their training methods, showing a preponderance of hypothetico-deductive reasoning and extensive elaborations of biomedical information. The PBL students have been shown to use hypothesis-driven reasoning—from the hypothesis to explain the patient data—while non-PBL students use mainly data-driven reasoning—from data towards the hypothesis In explaining clinical cases, PBL students produce extensive elaborations using detailed biomedical information, which is relatively absent from non-PBL students' explanations. However, these elaborations result in the generation of errors. PBL promotes the activation and elaboration of prior knowledge.

Patel and colleagues (Patel et al, 2001) also investigated the effects of non-PBL curricula on the use and integration of basic science and clinical knowledge and its relationship to reasoning in diagnostic explanation. The results showed that biomedical and clinical knowledge are not integrated and that very little biomedical information is used in routine problem-solving situations. There is significant use of expert-like data-driven strategies, however, in non-PBL students' explanations. The use of biomedical information increases when the clinical problems are complex; at the same time, hypothesis-driven strategies replace the data-driven strategies

Students from a PBL school integrated the two types of knowledge and in contrast to the non-PBL students, they spontaneously used biomedical information in solving even routine problems. We concluded that for students in the non-PBL curriculum, the clinical components of the problems are treated separately from the biomedical science components. The two components of the problem analysis seem to be viewed as serving different functions. However, when needed, the biomedical knowledge is utilized and seems to act as a “glue” that ties the two kinds of information together.

In the PBL curriculum, the integration of basic science and clinical knowledge is so tight that students appear unable to separate the two types of knowledge. As a result, PBL students generate unnecessarily elaborate explanations, leading to errors of reasoning. PBL seems to promote a type of learning in which basic biomedical knowledge becomes so tightly tied to specific clinical problem types that it becomes difficult to decouple this knowledge in context in order to transfer to a new situation (Anderson, Reder & Simon,1996; Holyoak,1984).

This outcome is consistent with how biomedical information is taught in the classroom in PBL schools, namely, by encouraging use of the hypothetico-deductive method, resulting in a predominantly backward-directed mode of reasoning. Elaborations are accompanied by a tendency to generate errors of scientific fact and flawed patterns of explanation, such as circular reasoning. Even though a student's explanation may be riddled with bugs and misconceptions, their harmful effects may be dependent on the direction of reasoning. If they reason forward, then they are likely to view their existing knowledge as adequate. In this case, misconceptions may be long-lasting and difficult to eradicate. If they reason backward, misconceptions might best be viewed as transient hypotheses which, in the light of experience, are either refuted or else modified to form the kernel of a more adequate explanation. Interestingly, differences in the patterns of reasoning acquired in both PBL and non-PBL medical curricula are found to be quite stable -- even after the students have completed medical school and are in residency training programs (Patel, Arocha, Lecissi, 2001: Patel & Kaufman, 2001).

Instruction that emphasizes decontextualized abstracted models of phenomena has not yielded much success in medicine or in other spheres of science education. It is widely believed that the amount of transfer will be a function of the overlap between the original domain of learning and the target domain. (Holyoak, 1984). PBL’s emphasis on real world problems represents a very good source of transfer to clinical situations. However, it is very challenging to create a problem set that most effectively embodies certain biomedical concepts while maximizing transfer. Knowledge that is overly contextualized can actually reduce transfer.

Technology-mediated Reasoning

All technologies mediate human performance. Technologies, whether they be computer-based or in some other form, transform the ways individuals and groups behave. They do not merely augment, enhance or expedite performance, although a given technology may do all of these things. The difference is not one of quantitative change, but one that is qualitative in nature. Technology, tools, and artifacts not only enhance people's ability to perform tasks but also change the way they perform tasks. In cognitive science, this ubiquitous phenomenon is called the representational effect, which refers to the phenomenon that different representations of a common abstract structure can generate dramatically different representational efficiencies, task complexities, and behavioral outcomes (Zhang & Norman, 1994).

Technology as External Representations

One approach to the study of how technology mediates thinking and reasoning is to consider technology as external representations (Zhang & Norman, 1994, 1995; Zhang, 1997). External representations are the knowledge and structure in the environment, as physical symbols, objects, or dimensions (e.g., written symbols, beads of abacuses, dimensions of a graph), and as external rules, constraints, or relations embedded in physical configurations (e.g., spatial relations of written digits, visual and spatial layouts of diagrams, physical constraints in abacuses). The information in external representations can be picked up, analyzed, and processed by perceptual systems alone, although the top-down participation of conceptual knowledge from internal representations can sometimes facilitate or inhibit the perceptual processes. External representations are more than inputs and stimuli to the internal mind. For many tasks, external representations are intrinsic components, without which the tasks either cease to exist or completely change in nature.

Diagrams, graphs, pictures, and information displays are typical external representations. They are used in many cognitive tasks such as problem solving, reasoning, and decision-making. In the studies of the relationship between mental images and external pictures, Chambers and Reisberg (1985; Reisberg, 1987) showed that external representations could give people access to knowledge and skills that are unavailable from internal representations. This advantage typically arises because internal representations are representations that are already interpreted and difficult to change, whereas external representations are subject to interpretations and thus can lead to different understandings under different conditions. In their studies of diagrammatic problem solving, Larkin & Simon (1987; Larkin, 1989) show that diagrammatic representations help reasoning and problem solving because they support operators that can recognize features easily and make inferences directly. In studies of logical reasoning with diagrams, Stenning and Oberlander (1995) demonstrated that diagrammatic representations such as Euler circles limit abstraction and thereby ease processing effort. It is well-known that different forms of graphic displays have different representational efficiencies for different tasks and can cause different cognitive behaviors. For example, Kleinmuntz and Schkade (1993) showed that different representations (graphs, tables, and lists) of the same information can dramatically change decision making strategies: with a tabular display people made one decision but with a graph display of the same information people made a different decision.

The Impact of Technology on Thinking in Medicine

The mediating role of technology can be evaluated at several levels of analysis. For example, electronic medical records alter the practice of individual clinicians in significant ways as discussed below. Changes to an information system substantially impacts organizational and institutional practices from research to billing to quality assurance. Even the introduction of patient-centered medical records early in the twentieth century necessitated changes in hospital architecture and considerably effected work practices in clinical settings. Salomon, Perkins, and Globerson (1987) introduce a useful distinction in considering the mediating role of technology on individual performance, the *effects with* technology and the *effects of* technology. The former is concerned with the changes in performance displayed by users while equipped with the technology. For example, when using an effective medical information system, physicians should be able to gather information more systematically and efficiently. In this capacity, medical information technologies may alleviate some of the cognitive load associated with a given task and permit them to focus on higher-order thinking skills, such as hypothesis generation and evaluation. The *effects of* technology refer to enduring changes in general cognitive capacities (knowledge and skills) as a consequence of interaction with a technology. For example, frequent use of information technologies may result in lasting changes in medical decision-making practices even in the absence of the system.

In several studiers involving mediating role of technology in clinical practice, Patel and her colleagues (Patel et al, 2000) observed the change of thinking and reasoning patterns caused by the change in methods of writing patient records: from paper records to electronic medical records (EMR). They found that before using EMR, physicians focus on exploration and discovery, use complex propositions, and tend to use data-driven reasoning. After using EMR, which has structured data, physicians focus on problem solving, use simple propositions, and tend to use problem-directed and hypothesis-driven reasoning. The change of behavior caused by the use of EMR remains when the physicians go back to paper records, showing the enduring effects of technology on human reasoning in medicine.

As the basis for many medical decisions, diagnostic reasoning requires the collection, understanding, and use of many types of patient information, such as history, lab results, symptoms, prescriptions, images, and so on. It is affected by not just the expertise of the clinicians but also by the way the information is acquired, stored, processed, and presented. If we consider clinicians as rational decision makers, the format of a display, as long as it contains the same information, should not affect the outcome of the reasoning and decision making process. But the formats of displays do affect many aspects of clinicians' task performance. Recently there are several studies on how different displays of information in EMR affect clinicians' behavior. Three major types of displays have been studied: source-based, time-based, and concept-based. Source-based displays organize medical data by the sources of the data, such as encounter notes, lab reports, medications, lab results, radiology imaging and report, physical exams, and so on. Time-based displays organize medical data as a temporal history of patient data. Concept-based displays organize medical data by clinically meaningful concepts or problems. In this case all data that are related to each specific problem are displayed together. For example, if a patient has symptoms such as coughing, chest pain, and fever, the lab results, imaging reports, prescriptions, assessments and plans are displayed together. In a study by Zeng et al. (2002), they found that different displays were good for different tasks. For example, source-based displays are good for clinicians to retrieve information for a specific test or procedure from a specific department, whereas concept-based displays are good for the search of information related to a specific disease.

With the rapid growth of computer-based information systems we are interacting more and more with computer-generated health information displays. To make these displays effectively and accurately generate the information that people need for informed reasoning, a good design of these displays is needed.

Conclusions and Future Directions

Investigations into the process of medical reasoning have been one area where advances in cognitive science have made significant contributions. In particular, reasoning in a medical context involving dense population and high degree of uncertainty (such as critical care environments), compounded with constraints imposed by resource availability, leads to increased use of heuristic strategies. The utility of heuristics lies in limiting the extent of purposeful search through data sets. By reducing redundancy they have substantial practical value. A significant part of a physician's cognitive effort is based on heuristic thinking. However, the use of heuristics introduces considerable bias in medical reasoning, often resulting in a number of conceptual and procedural errors. These include misconceptions about laws governing probability, instantiation of general rules to a specific patient at the point of care, prior probabilities and actions, as well as false validation. Much of physicians' reasoning is inductive with attached probability. Human thought is fallible and we cannot appreciate the fallibility of our thinking unless we draw on the understanding of how physicians' thinking processes operate in the real working environment.

Cognitive studies are increasingly moving towards investigations of "real-world" phenomena. The constraints of laboratory-based work prevent capturing the dynamics of real-world problems. This problem is particularly salient in high velocity critical care environments. In the best-case scenarios, this is creating the potential for great synergy between laboratory-based research and cognitive studies in the "wild". As discussed in this chapter, studies of thinking and reasoning in medicine, including a focus on medical errors and technology-mediated cognition, are increasingly paying attention to dimensions of medical work in clinical settings. The recent concern with understanding and reducing medical errors provides an opportunity for cognitive scientist to apply cognitive theories and methodologies to a pressing practical problem. A trend in health care, spurred partly by the advent of information technologies that foster communication, is the change in healthcare systems in that they are more and more multidisciplinary, collaborative and often span geographic regions. In addition, increasing costs of health care and rapid knowledge growth have also accelerated the trend towards collaboration of health care professionals to share knowledge and skills. Comprehensive patient care necessitates the communication of health-care providers in different medical domains, thereby optimizing the use of their expertise. Research on reasoning will need to continue to move towards a distributed model of cognition. This model will include a focus on both socially shared and technology-mediated reasoning.

Acknowledgment

This chapter is dedicated to the memory of late Yogesh C. Patel who devoted his life to the advancement of biomedical science. Through his deeds and words, he inspired us to devote ourselves to superior pursuits and to aspire to higher scientific standards.

References

Anderson, J. R., Reder, L. M., & Simon, H. A. (1996). Situated learning and education. *Educational Researcher, 25*(4), 5-11.

Arocha, J. F., & Patel, V. L. (1995). Construction-integration theory and clinical reasoning. In C. A. I. Weaver & S. Mannes (Eds.), *Discourse comprehension: Essays in honor of Walter Kintsch* (pp. 359-381).

Blois, M. S. (1990). Medicine and the nature of vertical reasoning. *New England Journal of Medicine, 318,* 847-851.

Bobrow, D. G. (Ed.). (1985). *Qualitative reasoning about physical systems* (1st MIT Press ed.). Cambridge, Mass.: MIT Press.

Boshuizen, H. P. A., & Schmidt, H. G. (1992). On the role of biomedical knowledge in clinical reasoning by experts, intermediates, and novices. *Cognitive Science, 16*(2), 153-184.

Brooks, L. R., Norman, G. R., & Allen, S. W. (1991). Role of specific similarity in a medical diagnostic task. *Journal of Experimental psychology: General, 120*(3), 278-287.

Chambers, D., & Reisberg, D. (1985). Can mental images be ambiguous? *Journal of Experimental Psychology: Human Perception & Performance, 11*(3), 317-328.

Chandrasekaran, B. (1994). The functional representation and causal process. In M. Yovitz (Ed.), *Advances in computing*. New York: Academic Press.

Chandrasekaran, B., Smith, J. W., & Sticklen, J. (1989). Deep models and their relation to diagnosis. *Artificial Intelligence in Medicine, 1,* 29-40.

Chapman, G. B., & Elstein, A. S. (2000). Cognitive processes and biases in medical decision making. In G. B. Chapman & A. Frank (Eds.), *Decision making in health care: Theory, psychology, and applications* (pp. 183-210).

Chi, M. T. H., Bassok, M., Lewis, M. W., Reiman, P., & Glaser, R. (1989). Self explanations: How students study and use examples in learning to solve problems. *Cognitive Science, 13,* 145-182.

Chi, M. T. H., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*, 121-152.

Clancey, W. J. (1985). Heuristic classification. *Artificial intelligence, 27*, 289-350.

Coiera, E. (2000). When conversation is better than computation. *Journal of the American Medical Informatics Association*, *7*(3), 277-286.

Dowie, J., & Elstein, A. S. (Eds.). (1988). *Professional judgment: A reader in clinical decision making*. Cambridge: Cambridge University Press.

Elstein, A. S., Kleinmuntz, B., Rabinowitz, M., McAuley, R., Murakami, J., Heckerling, P. S., et al. (1993). Diagnostic reasoning of high- and low-domain-knowledge clinicians: a reanalysis. *Medical Decision Making, 13*(1), 21-29.

Elstein, A. S., Shulman, L. S., & Sprafka, S. A. (1978). *Medical problem solving: An analysis of clinical reasoning.* Cambridge, MA: Harvard University Press.

Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: verbal reports as data* (Revised ed.). Cambridge, Mass: MIT Press.

Ericsson, K. A., & Smith, J. (1991). *Toward a general theory of expertise: Prospects and limits*. New York: Cambridge University Press.

Eshelman, L. (1988). MOLE: A knowledge acquisition tool for Cover-and-Differentiatiate systems. In S. C. Marcus (Ed.), *Automating Knowledge Acquistion for Expert Systems* (pp. 37-80). Boston: Kluwer.

Feltovich, P. J., Johnson, P. E., Moller, J. H., & Swanson, D. B. (1984). LCS: The role and development of medical knowledge in diagnostic expertise. In W. J. Clancey & E. H. Shortliffe (Eds.), *Readings in medical artificial intelligence: the first decade* (pp. 275-319). Reading, Mass: Addison-Wesley.

Flexner, A. (1910). *Medical education in the United States and Canada. A report to the Carnegie Foundation for the Advancement of Teaching*. Boston: Updyke.

Fordyce, J., Blank, F. S., Pekow, P., Smithline, H. A., Ritter, G., Gehlbach, S., et al. (2003). Errors in a busy emergency department. *Annals of Emergency Medicine, 42*(3), 324-333.

Frensch, P. A., & Funke, J. (1995). *Complex problem solving: The European perspective*. Hillsdale, NJ: Lawrence Erlbaum.

Frijda, N. H., & Elshout, J. J. (1979). Problem solving and thinking. In J. A. Michon, E. G. J. Eijkman & L. F. W. De Klerk (Eds.), *Handbook of psychonomics* (Vol. 2). Amsterdam: North-Holland.

Gaba, D. M. (1992). Dynamic decision-making in anesthesiology: Cognitive models and training approaches. In D. A. Evans & V. L. Patel (Eds.), *Advanced models of cognition for medical training and practice* (pp. 123-147). New York: Springer-Verlag.

Gentner, D., & Holyoak, K. J. (1997). Reasoning and learning by analogy: Introduction. *American Psychologist, 52*(1), 32-34.

Gorowitz, S., & McIntyre, A. (1978). Toward a theory of medical fallibility. *Journal of Medicine and Philosophy, 1*, 51-71.

Gorry, G. A. (1973). Computer-assisted clinical decision-making. *Methods of Information in Medicine*, *12*(1), 45-51.

Hardiman, P. T., Dufresne, R., & Mestre, J. P. (1989). The relation between problem categorization and problem solving among experts and novices. *Memory & Cognition, 17*(5), 627-638.

Hastie, R. (2001). Problems for judgment and decision making. *Annual Review of Psychology, 52*, 653-683.

Holland, J. H. (1986). *Induction : processes of inference, learning, and discovery*. Cambridge, Mass.: MIT Press.

Holyoak, K. J. (1985). The pragmatics of analogical transfer. *The Psychology of Learning and Motivation, 19*, 59-87.

Holyoak, K. J., & Thagard, P. (1997). The analogical mind. *American Psychologist*, *52*(1), 35-44.

Joseph, G. M., & Patel, V. L. (1990). Domain knowledge and hypothesis generation in diagnostic reasoning. *Medical Decision Making, 10*(1), 31-46.

Kassirer, J. P. (1989). Diagnostic reasoning. *Annals of Internal Medicine, 110*(11), 893-900.

Kassirer, J. P., & Kopelman, R. I. (1990). Diagnosis and the structure of memory. 2. Exemplars, scripts, and simulation. *Hospital Practice (Office Edition), 25*(11), 29-33, 36.

Klein, G. A. (1993). *Decision making in action : models and methods*. Norwood, N.J.: Ablex Pub.

Kleinmuntz, D. N., & Schkade, D. A. (1993). Information displays and decision processes. *Psychological Science, 4*(4), 221-227.

Kohn, L. T., Corrigan, J., & Donaldson, M. S. (1999). *To err is human: building a safer health system*. Washington, D.C: National Academy Press.

Kuhn, D. (1995). Scientific thinking and knowledge acquisition. *Monographs of the Society for Research in Child Development*, *60*(4), 152-157.

Kuipers, B. (1987). Qualitative simulation as causal explanation. *IEEE Transactions on Systems, Man, and Cybernetics, 17*, 432-444.

Kuipers, B., & Kassirer, J. P. (1984). Causal reasoning in medicine: Analysis of a protocol. *Cognitive Science, 8*(4), 363-385.

Larkin, J., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science, 11*(1), 65-99.

Leaper, D. J., Horrocks, J. C., Staniland, J. R., & De Dombal, F. T. (1972). Computer-assisted diagnosis of abdominal pain using "estimates" provided by clinicians. *British Medical Journal, 4*(836), 350-354.

Ledley, R. S., & Lusted, L. B. (1959). Reasnoning foundations of medical diagnosis. *Science, 130*, 9-21.

Leprohon, J., & Patel, V. L. (1995). Decision-making strategies for telephone triage in emergency medical services. *Med.Decis.Making, 15*(3), 240-253.

Lin, L., Isla, R., Doniz, K., Harkness, H., Vicente, K. J., & Doyle, D. J. (1998). Applying human factors to the design of medical equipment: patient-controlled analgesia. *Journal of Clinical Monitoring & Computing, 14*(4), 253-263.

Miller, R. A., Pople, H. E., & Myers, J. D. (1984). Internist-I, an experimental computer-based diagnostic for general internal medicine. In W. J. Clancey & E. H. Shortliffe (Eds.), *Readings in medical artificial intelligence: the first decade* (pp. xvi, 512 p). Reading, Mass: Addison-Wesley.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, Mass.: Harvard University Press.

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, N.J.: Prentice-Hall.

Norman, D. A. (1986). Cognitive engineering. In D. A. Norman & S. W. Draper (Eds.), *User centered system design: New perspectives on human-computer interaction* (pp. 31-61). Hillsdale, NJ: Lawrence Erlbaum Associates.

Norman, D. A. (1988). *The psychology of everyday things*. New York: Basic Books.

Norman, G. R., & Brooks, L. R. (1997). The non-analytical basis of clinical reasoning. *Advances in Health Sciences Education, 2*(2), 173-184.

Norman, G. R., Brooks, L. R., & Allen, S. W. (1989). Recall by expert medical practitioners and novices as a record of processing attention. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 15*(6), 1166-1174.

Osler, W. (1906). *Aequanimitas. With other addresses to medical students, nurses and practitioners of medicine*. Philadelphia: Blakiston's Son & Co.

Patel, V. L., & Arocha, J. F. (2001). The nature of constraints on collaborative decision making in health care settings. In E. Salas & G. Klein (Eds.), *Linking expertise and naturalistic decision making* (pp. 383-405). Mahwah, NJ: Lawrence Erlbaum Associates.

Patel, V. L., Arocha, J. F., & Kaufman, D. R. (2001). A primer on aspects of cognition for medical informatics. *J Am Med Inform Assoc, 8*(4), 324-343.

Patel, V.L., Arocha, J. Lecissi,M (2001) "Impact of Undergraduate Medical Training on Housestaff Problem Solving Performance, Implications for Health Education in Problem-based Curricula." *Journal of Dental Education*.65 (11) 1199-1218.

Patel, V. L., Evans, D. A., & Kaufman, D. R. (1989). Cognitive framework for doctor-patient communication. In D. A. Evans & V. L. Patel (Eds.), *Cognitive science in medicine: Biomedical modeling* (pp. 257-312). Cambridge, MA, US: MIT Press.

Patel, V. L., & Groen, G. J. (1986). Knowledge-based solution strategies in medical reasoning. *Cognitive Science, 10*, 91-116.

Patel, V. L., & Groen, G. J. (1991). The general and specific nature of medical expertise: A critical look. In K. A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise: Prospects and limits* (pp. 93-125). New York, NY: Cambridge University Press.

Patel, V. L., Groen, G. J., & Arocha, J. F. (1990). Medical expertise as a function of task difficulty. *Memory & Cognition, 18*(4), 394-406.

Patel, V. L., Groen, G. J., & Norman, G. R. (1993). Reasoning and instruction in medical curricula. *Cognition & Instruction, 10*(4), 335-378.

Patel, V. L., & Kaufman, D. R. (1994). Diagnostic reasoning and expertise. *Psychology of Learning and Motivation, 31*, 137-252.

Patel, V. L., & Kaufman, D. R. (1995). Clinical reasoning and biomedical knowledge: Implications for teaching. In J. Higgs & M. Jones (Eds.), *Clinical reasoning in the health professions* (pp. 117-128). Oxford: Butterworth Heinemenn.

Patel, V. L., & Kaufman, D. R. (2001, Feb 02). Medical education isn't just about solving problems. *The Chronicle of Higher Education,* p. B12.

Patel, V. L., Kaufman, D. R., Allen, V. G., Shortliffe, E. H., Cimino, J. J., & Greenes, R. A. (1999). Toward a framework for computer-mediated collaborative design in medical informatics. *Methods of Information in Medicine, 38*(3), 158-176.

Patel, V. L., Kaufman, D. R., & Arocha, J. F. (2002). Emerging paradigms of cognition in medical decision-making. *Journal of Biomedical Informatics, 35*, 52-75.

Patel, V. L., Kaufman, D. R., & Magder, S. A. (1991). Causal reasoning about complex physiological concepts in cardiovascular physiology by medical students. *International Journal of Science Education, 13*, 171-185.

Patel, V. L., Kushniruk, A. W., Yang, S., & Yale, J. F. (2000). Impact of a computer-based patient record system on data collection, knowledge organization, and reasoning. *Journal of the American Medical Informatics Association, 7*(6), 569-585.

Patel, V. L., & Ramoni, M. F. (1997). Cognitive models of directional inference in expert medical reasoning. In P. J. Feltovich & K. M. Ford (Eds.), *Expertise in context: Human and machine* (pp. 67-99): Cambridge, MA, US: The MIT Press.

Patil, R. S., Szolovits, P., & Schwartz, W. B. (1985). Causal understanding of patient illness in medical diagnosis. In J. A. Reggia & S. Tuhrim (Eds.), *Computer-assisted medical decision making* (Vol. 2, pp. 272-292). New York: Springer-Verlag.

Peirce, C. S. (1955). Abduction and induction. In C. S. Peirce & J. Buchler (Eds.), *Philosophical writings of Peirce* (pp. 150-156). New York: Dover Publications.

Perkins, D. N., & Simmons, R. (1988). An integrative model of misconceptions. *Review of Educational Research, 58*, 303-326.

Ramoni, M. F., Stefanelli, M., Magnani, L., & Barosi, G. (1992). An epistemological framework for medical knowledge based system. *IEEE Transactions on Systems, Man, and Cybernetics, 22*, 1361-1375.

Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive systems engineering*. New York: Wiley.

Reason, J. T. (1990). *Human error*. Cambridge, England: Cambridge University Press.

Reisberg, D. (1987). External representations and the advantages of externalizing one's thoughts. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Rimoldi, H. J. A. (1961). The test of diagnostic skills. *Journal of Medical Education, 36*, 73-79.

Rogoff, B., & Lave, J. (1984). *Everyday cognition : its development in social context*. Cambridge, Mass.: Harvard University Press.

Salas, E., & Klein, G. A. (2001). *Linking expertise and naturalistic decision making*. Mahwah, N.J.: L. Erlbaum Associates.

Salomon, G., & Perkins, D. N. (1989). Rocky roads to transfer: Rethinking mechanisms of a neglected phenomenon. *Educational Psychologist, 24*, 113-142.

Schaffner, K. F. (1986). Exemplar reasoning about biological models and diseases: A relation between the philosophy of medicine and philosophy of science. *Journal of Medicine & Philosophy, 11*, 63-80.

Schauble, L. (1996). The development of scientific reasoning in knowledge-rich contexts. *Developmental Psychology, 32*(1), 102-119.

Schmidt, H. G., & Boshuizen, H. P. A. (1993). On the origin of intermediate effects in clinical case recall. *Memory & Cognition, 21*, 338-351.

Shortliffe, E. H. (1976). *Computer-based medical consultations, MYCIN*. New York: Elsevier.

Simon, D. P., & Simon, H. A. (1978). Individual differences in solving physics problems. In R. Siegler (Ed.), *Children's thinking: What developes?* Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.

Stefanelli, M., & Ramoni, M. F. (1992). Epistemological constraints on medical knowledge-based systems. In D. A. Evans & V. L. Patel (Eds.), *Advanced models of cognition for medical training and practice* (Vol. 97, pp. 3-20). Heidelberg, Germany: Springer-Verlag.

Stenning, K., & Oberlander, J. (1994). A cognitive theory of graphical and linguistic reasoning: Logic and implementation. *Cognitive Science, 19*, 97-140.

Strauss, S., & Stavy, R. (1982). *U-shaped behavioral growth*. New York: Academic Press.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12, 257-85.

Szolovits, P. (Ed.). (1982). *Artificial intelligence in medicine* (Vol. 51). Boulder, Colo.: Published by Westview Press for the American Association for the Advancement of Science.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124-1131.

Vicente, K. J., & Rasmussen, J. (1990). The ecology of human-machine systems. II: Mediating "direct perception" in complex work domains. *Ecological Psychology, 2*, 207-250.

White, B. Y., & Frederiksen, J. R. (1990). Causal model progressions as a foundation for intelligent learning environments. In W. J. Clancey & E. Soloway (Eds.), *Artificial intelligence and learning environments Special issues of "Artificial Intelligence: An International Journal"* (pp. 99-157).

Woods, D. D. (1993). Process-tracing methods for the study of cognition outside of the experimental psychology laboratory. In G. A. Klein & J. Orasanu (Eds.), *Decision making in action: Models and methods* (pp. 228-251). Norwood, NJ: Ablex.

Woods, D. D., & Cook, R. (1998). *A tale of two stories, contrasting views of patient safety*: National Health Care Safety Council of the National Patient Safety Foundation at the AMA.

Zeng, Q., Cimino, J. J., & Zou, K. H. (2002). Providing concept-oriented views for clinical data using a knowledge-based system: an evaluation. *Journal of the American Medical Informatics Association, 9*(3), 294-305.

Zhang, J. (1996). A representational analysis of relational information displays. *International Journal of Human-Computer Studies, 45*(1), 59-74.

Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive Science, 21*(2), 179-217.

Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science, 18*, 87-122.

Zhang J., Patel VL, Johnson TR, Shortliffe EH (under review, 2003) A Cognitive Taxonomy of Medical Errors, Journal of Biomedical Informatics.