

Vimla L. Patel and David R. Kaufman

After reading this chapter, you should know the answers to these questions:

- How can cognitive science theory meaningfully inform and shape design, development and assessment of healthcare information systems?
- What are some of the ways in which cognitive science differs from behavioral science?
- What are some of the ways in which we can characterize the structure of knowledge?
- What are the basic HCI and cognitive science methods that are useful for healthcare information system evaluation and design?
- What are some of the dimensions of difference between experts and novices?
- What are the attributes of system usability?
- What are the gulfs of execution and evaluation? What role do these considerations play in system design?
- Why is it important to consider cognition and human factors in dealing with issues of patient safety?

---

## 4.1 Introduction

Enormous advances in health information technologies and more generally, in computing over the course of the past two decades have begun to permeate diverse facets of clinical practice. The rapid pace of technological developments such as the Internet, wireless technologies, and hand-held devices, in the last decade affords significant opportunities for supporting, enhancing and extending user experiences, interactions and communications (Rogers 2004). These advances coupled with a growing computer literacy among healthcare professionals afford the potential for great improvement in healthcare. Yet many observers note that the healthcare system is slow to understand information technology and effectively incorporate it into the work environment (Shortliffe and Blois 2001). Innovative technologies often produce profound cultural, social, and cognitive changes. These transformations necessitate adaptation at many different levels of aggregation from the individual to the larger institution, sometimes causing disruptions of workflow and user dissatisfaction.

Similar to other complex domains, biomedical information systems embody ideals in design that often do not readily yield practical solutions in implementation. As computer-based systems infiltrate clinical practice and settings, the consequences often can be felt through all levels of the organization. This impact can have deleterious effects resulting in systemic inefficiencies and suboptimal practice, which can lead to frustrated

---

V.L. Patel, PhD, DSc (✉)  
Department of Center for Cognitive Studies in  
Medicine and Public Health, The New York Academy  
of Medicine, 1216 Fifth Avenue,  
New York, NY 10029, USA  
e-mail: vpatel@nyam.org, vimla.patel@uth.tmc.edu

D.R. Kaufman, PhD  
Biomedical Informatics, Arizona State University,  
13212 East Shea Blvd., Scottsdale, AZ 85259  
e-mail: kaufman@dbmi.columbia.edu

healthcare practitioners, unnecessary delays in healthcare delivery, and even adverse events (Lin et al. 1998; Weinger and Slagle 2001). In the best-case scenario, mastery of the system necessitates an individual and collective learning curve yielding incremental improvements in performance and satisfaction. In the worst-case scenario, clinicians may revolt and the hospital may be forced to pull the plug on an expensive new technology. How can we manage change? How can we introduce systems that are designed to be more intuitive and also implemented to be coherent with everyday practice?

### 4.1.1 Introducing Cognitive Science

**Cognitive science** is a multidisciplinary domain of inquiry devoted to the study of cognition and its role in intelligent agency. The primary disciplines include cognitive psychology, artificial intelligence, neuroscience, linguistics, anthropology, and philosophy. From the perspective of informatics, cognitive science can provide a framework for the analysis and modeling of complex human performance in technology-mediated settings. Cognitive science incorporates basic science research focusing on fundamental aspects of cognition (e.g., attention, memory, reasoning, early language acquisition) as well as applied research. Applied cognitive research is focally concerned with the development and evaluation of useful and usable cognitive artifacts. **Cognitive artifacts** are human-made materials, devices, and systems that extend people's abilities in perceiving objects, encoding and retrieving information from memory, and problem-solving (Gillan and Schvaneveldt 1999). In this regard, applied cognitive research is closely aligned with the disciplines of **human-computer interaction** (HCI) and **human factors**. It also has a close affiliation with educational research. In everyday life, we interact with cognitive artifacts to receive and/or manipulate information so as to alter our thinking processes and offload effort-intensive cognitive activity to the external world, thereby reducing mental workload.

The past couple of decades have produced a cumulative body of experiential and practical

knowledge about system design and implementation that can guide future initiatives. This practical knowledge embodies the need for sensible and intuitive user interfaces, an understanding of workflow, and the ways in which systems impact individual and team performance. However, experiential knowledge in the form of anecdotes and case studies is inadequate for producing robust generalizations or sound design and implementation principles. There is a need for a theoretical foundation. Biomedical informatics is more than the thin intersection of biomedicine and computing (Patel and Kaufman 1998). There is a growing role for the social sciences, including the cognitive and behavioral sciences, in biomedical informatics, particularly as they pertain to human-computer interaction and other areas such as **information retrieval** and **decision support**. In this chapter, we focus on the foundational role of cognitive science in biomedical informatics research and practice. Theories and methods from the cognitive sciences can illuminate different facets of design and implementation of information and knowledge-based systems. They can also play a larger role in characterizing and enhancing human performance on a wide range of tasks involving clinicians, patients and healthy consumers of biomedical information. These tasks may include developing training programs and devising measures to reduce errors or increase efficiency. In this respect, cognitive science represents one of the component basic sciences of biomedical informatics (Shortliffe and Blois 2001).

### 4.1.2 Cognitive Science and Biomedical Informatics

How can cognitive science theory meaningfully inform and shape design, development and assessment of health-care information systems? Cognitive science provides insight into principles of system usability and *learnability*, the mediating role of technology in clinical performance, the process of medical judgment and decision-making, the training of healthcare professionals, patients and health consumers, and

**Table 4.1** Correspondences between cognitive science, medical cognition and applied cognitive research in medical informatics

Cognitive Science	Medical Cognition	Biomedical Informatics
Knowledge organization and human memory	Organization of clinical and basic science knowledge	Development and use of medical knowledge bases
Problem solving, Heuristics/reasoning strategies	Medical problem solving and decision making	Medical artificial intelligence/decision support systems/medical errors
Perception/attention	Radiologic and dermatologic diagnosis	Medical imaging systems
Text comprehension	Learning from medical texts	Information retrieval/digital libraries/health literacy
Conversational analysis	Medical discourse	Medical natural language processing
Distributed cognition	Collaborative practice and research in health care	Computer-based provider order entry systems
Coordination of theory and evidence	Diagnostic and therapeutic reasoning	Evidence-based clinical guidelines
Diagrammatic reasoning	Perceptual processing of patient data displays	Biomedical information visualization

the design of a safer workplace. The central argument is that it can inform our understanding of human performance in technology-rich health-care environments (Carayon 2007).

Precisely how will cognitive science theory and methods make such a significant contribution towards these important objectives? The translation of research findings from one discipline into practical concerns that can be applied to another is rarely a straight-forward process (Rogers 2004). Furthermore, even when scientific knowledge is highly relevant in principle, making that knowledge effective in a design context can be a significant challenge. In this chapter, we discuss (a) basic cognitive science research and theories that provide a foundation for understanding the underlying mechanisms guiding human performance (e.g., findings pertaining to the structure of human memory), and (b) research in the areas of **medical errors** and **patient safety** as they interact with health information technology),

As illustrated in Table 4.1, there are correspondences between basic cognitive science research, medical cognition and cognitive research in biomedical informatics along several dimensions. For example, theories of human memory and knowledge organization lend themselves to characterizations of expert clinical knowledge that can then be contrasted with representation

of such knowledge in clinical systems. Similarly, research in text comprehension has provided a theoretical framework for research in understanding biomedical texts. This in turn has influenced applied cognitive research on information retrieval (Chap. 21) from biomedical knowledge sources and research on **health literacy**. Similarly, theories of problem solving and reasoning can be used to understand the processes and knowledge associated with diagnostic and therapeutic reasoning. This understanding provides a basis for developing biomedical artificial intelligence and decision support systems.

In this chapter, we demonstrate that cognitive research, theories and methods can contribute to applications in informatics in a number of ways including: (1) seed *basic research findings* that can illuminate dimensions of design (e.g., attention and memory, aspects of the visual system), (2) provide an *explanatory vocabulary* for characterizing how individuals process and communicate health information (e.g., various studies of medical cognition pertaining to doctor-patient interaction), (3) present an *analytic framework* for identifying problems and modeling certain kinds of user interactions, (4) characterize the relationship between health information technology, human factors and patient safety, (5) provide *rich descriptive accounts* of clinicians employing technologies in

the context of work, and (6) furnish a generative approach for novel designs and productive applied research programs in informatics (e.g., intervention strategies for supporting low literacy populations in health information seeking).

Since the last edition of this text, there has been a significant growth in cognitive research in biomedical informatics. We conducted an informal comparison of studies across three leading informatics journals, *Journal of Biomedical Informatics*, *Journal of the American Medical Informatics Association* and the *International Journal of Medical Informatics* of two time periods over the last decade, the first being 2001–2005 and the second 2006–2010. A keyword search of ten common terms (e.g., cognition, usability testing and human factors) found an increase of almost 70 % in the last 5 years over the previous 5 years. Although this doesn't constitute a rigorous systematic analysis, it is suggestive of a strong growth of cognitive research in informatics.

The social sciences are constituted by multiple frameworks and approaches. **Behaviorism** constitutes a framework for analyzing and modifying behavior. It is an approach that has had an enormous influence on the social sciences for most of the twentieth Century. Cognitive science partially emerged as a response to the limitations of behaviorism. The next section of the chapter contains a brief history of the cognitive and behavioral sciences that emphasizes the points of difference between the two approaches. It also serves to introduce basic concepts in the study of cognition.

---

## 4.2 Cognitive Science: The Emergence of an Explanatory Framework

Cognitive science is, of course, not really a new discipline, but recognition of a fundamental set of common concerns shared by the disciplines of psychology, computer science, linguistics, economics, epistemology, and the social sciences generally. All of these disciplines are concerned with information processing systems, and all of them

are concerned with systems that are adaptive—that are what they are from being ground between the nether millstone of their physiology or hardware, as the case may be, and the upper millstone of a complex environment in which they exist. Herbert A. Simon (1980. P 33) (H. A. Simon 1980)

In this section, we sketch a brief history of the emergence of cognitive science in view to differentiate it with competing theoretical frameworks in the social sciences. The section also serves to introduce core concepts that constitute an explanatory framework for cognitive science.

Behaviorism is the conceptual framework underlying a particular science of behavior (Zuriff 1985). This framework dominated experimental and applied psychology as well as the social sciences for the better part of the twentieth century (Bechtel et al. 1998). Behaviorism represented an attempt to develop an objective, empirically based science of behavior and more specifically, learning. Empiricism is the view that experience is the only source of knowledge (Hilgard and Bower 1975). Behaviorism endeavored to build a comprehensive framework of scientific inquiry around the experimental analysis of observable behavior. Behaviorists eschewed the study of thinking as an unacceptable psychological method because it was inherently subjective, error prone, and could not be subjected to empirical validation. Similarly, hypothetical constructs (e.g., mental processes as mechanisms in a theory) were discouraged. All constructs had to be specified in terms of operational definitions so they could be manipulated, measured and quantified for empirical investigation (Weinger and Slagle 2001). Radical behaviorism as espoused by B.F. Skinner proposed that behavioral events may be understood and analyzed entirely in relation to past and present environment and evolutionary history without any reference to internal states (Baum 2011).

Behavioral theories of learning emphasized the correspondence between environmental stimuli and the responses emitted. These studies generally attempted to characterize the changing relationship between stimulus and response under different reinforcement and punishment contingencies. For example, a behavior that was followed by a satisfying state of affairs is more

likely to increase the frequency of the act. According to behavior theories, knowledge is nothing more than sum of an individual's learning history and transformations of mental states play no part in the learning process.

For reasons that go beyond the scope of this chapter, classical behavioral theories have been largely discredited as a comprehensive unifying theory of behavior. However, behaviorism continues to provide a theoretical and methodological foundation in a wide range of social science disciplines. For example, behaviorist tenets continue to play a central role in public health research. In particular, health behavior research places an emphasis on antecedent variables and environmental contingencies that serve to sustain unhealthy behaviors such as smoking (Sussman 2001). Around 1950, there was an increasing dissatisfaction with the limitations and methodological constraints (e.g., the disavowal of the unobserved such as mental states) of behaviorism. In addition, developments in logic, information theory, cybernetics, and perhaps most importantly the advent of the digital computer, aroused substantial interest in "information processing" (Gardner 1985).


Newell and Simon (1972) date the beginning of the "cognitive revolution" to the year 1956. They cite Bruner, Goodnow and Austin's "Study of Thinking," George Miller's influential journal publication "The magic number seven" in psychology, Noam Chomsky's writings on syntactic grammars in linguistics (see Chap. 8), and their own logic theorist program in computer science as the pivotal works. Cognitive scientists placed "thought" and "mental processes" at the center of their explanatory framework.

The "computer metaphor" provided a framework for the study of human cognition as the manipulation of "symbolic structures." It also provided the foundation for a model of memory, which was a prerequisite for an information processing theory (Atkinson and Shiffrin 1968). The implementation of models of human performance as computer programs provided a measure of objectivity and a *sufficiency test* of a theory and also serves to increase the objectivity of the study of mental processes (Estes 1975).

Arguably, the most significant landmark publication in the nascent field of cognitive science is Newell and Simon's "Human Problem Solving" (Newell and Simon 1972). This was the culmination of more than 15 years of work on problem solving and research in artificial intelligence. It was a mature thesis that described a theoretical framework, extended a language for the study of cognition, and introduced protocol-analytic methods that have become ubiquitous in the study of high-level cognition. It laid the foundation for the formal investigation of symbolic-information processing (more specifically, problem solving). The development of models of human information processing also provided a foundation for the discipline of human-computer interaction and the first formal methods of analysis (Card et al. 1983).

The early investigations of problem solving focused primarily on investigations of experimentally contrived or toy-world tasks such as elementary deductive logic, the Tower of Hanoi, illustrated in Fig. 4.1, and mathematical word problems (Greeno and Simon 1988). These tasks required very little background knowledge and were well structured, in the sense that all the variables necessary for solving the problem were present in the problem statement. These tasks allowed for a complete description of the task environment; a step-by-step description of the sequential behavior of the subjects' performance; and the modeling of subjects' cognitive and overt behavior in the form of a computer simulation. The Tower of Hanoi, in particular, served as an important test bed for the development of an explanatory vocabulary and framework for analyzing problem solving behavior.

The Tower of Hanoi (TOH) is a relatively straight-forward task that consists of three pegs (A, B, and C) and three or more disks that vary in size. The goal is to move the three disks from peg A to peg C one at a time with the constraint that a larger disk can never rest on a smaller one. Problem solving can be construed as *search* in a **problem space**. A problem space has an *initial state*, a *goal state*, and a *set of operators*. Operators are any moves that transform a given state to a successor state. For example, the first



**Fig. 4.1** Tower of Hanoi task illustrating a start state and a goal state

move could be to move the small disk to peg B or peg C. In a three-disk TOH, there are a total of 27 possible states representing the complete problem space. TOH has  $3^n$  states where  $n$  is the number of disks. The minimum number of moves necessary to solve a TOH is  $2^{n-1}$ . Problem solvers will typically maintain only a small set of states at a time.

The search process involves finding a solution strategy that will minimize the number of steps. The metaphor of movement through a problem space provides a means for understanding how an individual can sequentially address the challenges they confront at each stage of a problem and the actions that ensue. We can characterize the problem-solving behavior of the subject at a local level in terms of state transitions or at a more global level in terms of *strategies*. For example, *means ends analysis* is a commonly used strategy for reducing the difference between the start state and goal state. For instance, moving all but the largest disk from peg A to peg B is an interim goal associated with such a strategy. Although TOH bears little resemblance to the tasks performed by either clinicians or patients, the example illustrates the process of analyzing task demands and task performance in human subjects.

The most common method of data analysis is known as **protocol analysis**<sup>1</sup> (Newell and Simon 1972). Protocol analysis refers to a class of

techniques for representing verbal **think-aloud protocols** (Greeno and Simon 1988). Think aloud protocols are the most common source of data used in studies of problem solving. In these studies, subjects are instructed to verbalize their thoughts as they perform a particular experimental task. Ericsson and Simon (1993) specify the conditions under which verbal reports are acceptable as legitimate data. For example, retrospective think-aloud protocols are viewed as somewhat suspect because the subject has had the opportunity to reconstruct the information in memory and the verbal reports are inevitably distorted. Think aloud protocols recorded in concert with observable behavioral data such as a subject's actions provide a rich source of evidence to characterize cognitive processes.

Cognitive psychologists and linguists have investigated the processes and properties of language and memory in adults and children for many decades. Early research focused on basic laboratory studies of list learning or processing of words and sentences (as in a sentence completion task) (Anderson 1983). Beginning in the early 1970s, van Dijk and Kintsch (1983) developed an influential method of analyzing the process of **text comprehension** based on the realization that text can be described at multiple levels of realization from surface codes (e.g., words and syntax) to deeper level of semantics. Comprehension refers to cognitive processes associated with understanding or deriving meaning from text, conversation, or other informational resources. 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

<sup>1</sup>The term protocol refers to that which is produced by a subject during testing (e.g., a verbal record). It differs from the more common use of protocol as defining a code or set of procedures governing behavior or a situation.

such processes, which is, the mental representation of the text, essentially what people have understood.

Comprehension often precedes problem solving and decision making, but 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 more important differences in medical problem solving and decision making arise from differences in knowledge and comprehension. Furthermore, many of the problems associated with decision making are the result of either lack of knowledge or failure to understand the information appropriately.

The early investigations provided a well-constrained artificial environment for the development of the basic methods and principles of problem solving. They also provide a rich explanatory vocabulary (e.g., problem space), but were not fully adequate in accounting for cognition in knowledge-rich domains of greater complexity and involving uncertainty. In the mid to late 1970s, there was a shift in research to complex “real-life” knowledge-based domains of enquiry (Greeno and Simon 1988). Problem-solving research was studying performance in domains such as physics (1980), medical diagnoses (Elstein et al. 1978) and architecture (Akin 1982). Similarly the study of text comprehension shifted from research on simple stories to technical and scientific texts in a range of domains including medicine. This paralleled a similar change in artificial intelligence research from “toy programs” to addressing “real-world” problems and the development of expert systems (Clancey and Shortliffe 1984). The shift to real-world problems in cognitive science was spearheaded by research exploring the nature of expertise. Most of the early investigations on expertise involved laboratory experiments. However, the shift to knowledge-intensive domains provided a theoretical and methodological foundation to conduct both basic and applied research in real-world settings such as the workplace (Vicente 1999) and the classroom (Bruer 1993). These areas of application provided a fertile test bed for assessing and extending the cognitive science framework.

In recent years, the conventional information-processing approach has come under criticism for its narrow focus on the rational/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 solitary individual to viewing cognition as distributed across groups, cultures, and artifacts. This claim has significant implications for the study of collaborative endeavors and human-computer interaction. We explore the concepts underlying *distributed cognition* in greater detail in a subsequent section.

---

### 4.3 Human Information Processing

It is well known that product design often fails to adequately consider cognitive and physiological constraints and imposes an unnecessary burden on task performance (Preece et al. 2007). Fortunately, advances in theory and methods provide us with greater insight into designing systems for the human condition.

Cognitive science serves as a basic science 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 of 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 clinician’s hypothesis about a patient’s condition after noticing an abnormal gait as he entered the clinic. These are likely to elicit further inferences about the patient’s underlying condition and may direct the physician’s information-gathering strategies and contribute to an evolving problem representation.

Two interdependent dimensions by which we can characterize cognitive systems are: (1) architectural theories that endeavor to provide a unified theory for all aspects of cognition and

(2) distinction the different kinds of knowledge necessary to attain competency in a given domain. Individuals differ substantially in terms of their knowledge, experiences, and endowed capabilities. The architectural approach capitalizes on the fact that 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 lifespan. In regards to the lifespan issue, there is a growing body of literature on cognitive aging and how aspects of the cognitive system such as attention, memory, vision and motor skills change as a function of aging (Fisk et al. 2009). This basic science research is of growing importance to informatics as we seek to develop e-health applications for seniors, many of whom suffer from chronic health conditions such as arthritis and diabetes. A graphical user interface or more generally, a website designed for younger adults may not be suitable for older adults.

Differences in knowledge organization are a central focus of research into the nature of expertise. In medicine, the expert-novice paradigm has contributed to our understanding of the nature of medical expertise and skilled clinical performance.

### 4.3.1 Cognitive Architectures and Human Memory Systems

Fundamental research in perception, cognition, and psychomotor skills over the course of the last 50 years has provided a foundation for design principles in human factors and human-computer interaction. Although cognitive guidelines have made significant inroads in the design commu-

nity, there remains a significant gap in applying basic cognitive research (Gillan and Schvaneveldt 1999). Designers routinely violate basic assumptions about the human cognitive system. There are invariably challenges in applying basic research and theory to applications. A more human-centered design and cognitive research can instrumentally contribute to such an endeavor (Zhang et al. 2004).

Over the course of the last 25 years, there have been several attempts to develop a unified theory of cognition. The goal of such a theory is to provide a single set of mechanisms for all cognitive behaviors from motor skills, language, memory, to decision making, problem solving and comprehension (Newell 1990). Such a theory provides a means to put together a voluminous and seemingly disparate body of human experimental data into a coherent form. Cognitive architecture represents unifying theories of cognition that are embodied in large-scale computer simulation programs. Although there is much plasticity evidenced in human behavior, cognitive processes are bound by biological and physical constraints. Cognitive architectures specify functional rather than biological constraints on human behavior (e.g., limitations on working memory). These constraints reflect the information-processing capacities and limitations of the human cognitive system. Architectural systems embody a relatively fixed permanent structure that is (more or less) characteristic of all humans and doesn't substantially vary over an individual's lifetime. It represents a scientific hypothesis about those aspects of human cognition that are relatively constant over time and independent of task (Carroll 2003). Cognitive architectures also play a role in providing blueprints for building future intelligent systems that embody a broad range of capabilities similar to those of humans (Duch et al. 2008).

Cognitive architectures include short-term and long-term memories that store content about an individual's beliefs, goals, and knowledge, the representation of elements that are contained in these memories as well as their organization into larger-scale structures (Langley et al. 2009). An extended discussion of architectural theories and



systems is beyond the scope of this chapter. However, we employ the architectural frame of reference to introduce some basic distinctions in memory systems. Human memory is typically divided into at least two structures: **long-term memory** and **short-term/working memory**. Working memory is an emergent property of interaction with the environment. 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., text comprehension). The information maintained in working memory includes stimuli from the environment (e.g., words on a display) and knowledge activated from long-term memory. In theory, LTM is infinite, whereas WM is limited to five to ten “chunks” of information. A chunk is any stimulus or patterns of stimuli that has become familiar from repeated exposure and is subsequently stored in memory as a single unit (Larkin et al. 1980). Problems impose a varying **cognitive load** on working memory. This refers to an excess of information that competes for few cognitive resources, creating a burden on working memory (Chandler and Sweller 1991). For example, maintaining a seven-digit phone number in WM is not very difficult. However, to maintain a phone number while engaging in conversation is nearly impossible for most people. Multi-tasking is one factor that contributes to cognitive load. The structure of the task environment, for example, a crowded computer display is another contributor. High velocity/high workload clinical environments such as intensive care units also impose cognitive loads on clinicians carrying out task.

### 4.3.2 The Organization of Knowledge

Architectural theories specify the structure and mechanisms of memory systems, whereas theories of knowledge organization focus on the content. There are several ways to characterize the kinds of knowledge that reside in LTM and that support decisions and actions. Cognitive

psychology has furnished a range of domain-independent constructs that account for the variability of mental representations needed to engage the external world.

A central tenet of cognitive science is that humans actively construct and interpret information from their environment. Given that environmental stimuli can take a multitude of forms (e.g., written text, speech, music, images, etc.), the cognitive system needs to be attuned to different representational types to capture the essence of these inputs. For example, we process written text differently than we do mathematical equations. The power of cognition is reflected in the ability to form abstractions - to represent perceptions, experiences and thoughts in some medium other than that in which they have occurred without extraneous or irrelevant information (Norman 1993). Representations enable us to remember, reconstruct, and transform events, objects, images, and conversations absent in space and time from our initial experience of the phenomena. Representations reflect states of knowledge.

**Propositions** are a form of natural language representation that captures the essence of an idea (i.e., semantics) or concept without explicit reference to linguistic content. For example, “hello”, “hey”, and “what’s happening” can typically be interpreted as a greeting containing identical propositional content even though the literal semantics of the phrases may differ. These ideas are expressed as language and translated into speech or text when we talk or write. Similarly, we recover the propositional structure when we read or listen to verbal information. Numerous psychological experiments have demonstrated that people recover the gist of a text or spoken communication (i.e., propositional structure) not the specific words (Anderson 1985; van Dijk and Kintsch 1983). Studies have also shown the individuals at different levels of expertise will differentially represent a text (Patel and Kaufman 1998). For example, experts are more likely to selectively encode relevant propositional information that will inform a decision. On the other hand, non-experts will often remember more information, but much of the recalled information may not be relevant to the decision (Patel

**Fig. 4.2** Propositional analysis of a think-aloud protocol of a primary care physician

**1. 43-year-old white female who developed diarrhea after a brief period of 2 days of GI upset.**

1.1	female	ATT: Age (old); DEG: 43 year; ATT: white
1.2	develop	PAT: [she]; THM: diarrhea; TNS: past
1.3	period	ATT: brief; DUR: 2 days; THM: 1.4
1.4	upset	LOC: GI
1.5	TEM:ORD	[1.3], [1.2]

and Groen 1991a, b). Propositional representations constitute an important construct in theories of comprehension and are discussed later in this chapter.

Propositional knowledge can be expressed using a predicate calculus formalism or as a semantic network. The predicate calculus representation is illustrated below. A subject's response, as given on Fig. 4.2, is divided into sentences or segments and sequentially analyzed. The formalism includes a head element of a segment and a series of arguments. For example in proposition 1.1, the focus is on a female who has the attributes of being 43 years of age and white. The TEM:ORD or temporal order relation indicates that the events of 1.3 (GI upset) precede the event of 1.2 (diarrhea). The formalism is informed by an elaborate propositional language (Frederiksen 1975) and was first applied to the medical domain by Patel and her colleagues (Patel et al. 1986). The method provides us with a detailed way to characterize the information subjects understood from reading a text, based on their summary or explanations.

Kintsch (1998) theorized that comprehension involves an interaction between what the text conveys and the schemata in long-term memory. Comprehension occurs when the reader uses prior knowledge to process the incoming information presented in the text. The text information is called the *textbase* (the propositional content of the text). For instance, in medicine the textbase could consist of the representation of a patient problem as written in a patient chart. The situation model is constituted by the textbase representation plus the domain-specific and everyday knowledge that the reader uses to derive a broader meaning from the text. In medicine, the situation model would enable a physician to draw inferences from a patient's history leading to a diagnosis, therapeutic plan or prognosis (Patel

and Groen 1991a, b). This situation model is typically derived from the general knowledge and specific knowledge acquired through medical teaching, readings (e.g., theories and findings from biomedical research), clinical practice (e.g., knowledge of associations between clinical findings and specific diseases, knowledge of medications or treatment procedures that have worked in the past) and the textbase representation. Like other forms of knowledge representation, the situation model is used to "fit in" the incoming information (e.g., text, perception of the patient). Since the knowledge in LTM differs among physicians, the resulting situation model generated by any two physicians is likely to differ as well. Theories and methods of text comprehension have been widely used in the study of medical cognition and have been instrumental in characterizing the process of guideline development and interpretation (Arocha et al. 2005).

**Schemata** represent higher-level knowledge structures. They can be construed as data structures for representing categories of concepts stored in memory (e.g., fruits, chairs, geometric shapes, and thyroid conditions). There are schemata for concepts underlying situations, events, sequences of actions and so forth. To process information with the use of a schema is to determine which model best fits the incoming information. Schemata have constants (all birds have wings) and variables (chairs can have between one and four legs). The variables may have associated default values (e.g., birds fly) that represent the prototypical circumstance.

When a person interprets information, the schema serves as a "filter" for distinguishing relevant and irrelevant information. Schemata can be considered as generic knowledge structures that contain slots for particular kinds of propositions. For instance, a schema for myocardial infarction may contain the findings of "chest


pain,” “sweating,” “shortness of breath,” but not the finding of “goiter,” which is part of the schema for thyroid disease.

The schematic and propositional representations reflect abstractions and don’t necessarily preserve literal information about the external world. Imagine that you are having a conversation at the office about how to rearrange the furniture in your living room. To engage in such a conversation, one needs to be able to construct images of the objects and their spatial arrangement in the room. **Mental images** are a form of internal representation that captures perceptual information recovered from the environment. There is compelling psychological and neuropsychological evidence to suggest that mental images constitute a distinct form of mental representation (Bartolomeo 2008) Images play a particularly important role in domains of visual diagnosis such as dermatology and radiology.

**Mental models** are an analogue-based construct for describing how individuals form internal models of systems. Mental models are designed to answer questions such as “how does it work?” or “what will happen if I take the following action?” “Analogy” suggests that the representation explicitly shares the structure of the world it represents (e.g., a set of connected visual images of a partial road map from your home to your work destination). This is in contrast to an abstraction-based form such as propositions or schemas in which the mental structure consists of either the gist, an abstraction, or summary representation. However, like other forms of mental representation, mental models are always incomplete, imperfect and subject to the processing limitations of the cognitive system. Mental models can be derived from perception, language or from one’s imagination (Payne 2003). *Running* of a model corresponds to a process of mental simulation to generate possible future states of a system from observed or hypothetical state. For example, when one initiates a Google Search, one may reasonably anticipate that system will return a list of relevant (and less than relevant) websites that correspond to the query. Mental models are a particularly useful construct in understanding human-computer interaction.

An individual’s mental models provide predictive and explanatory capabilities of the function of a physical system. More often the construct has been used to characterize models that have a spatial and temporal context, as is the case in reasoning about the behavior of electrical circuits (White and Frederiksen 1990). The model can be used to simulate a process (e.g., predict the effects of network interruptions on getting cash from an ATM machine). Kaufman, Patel and Magder (1996) characterized clinicians’ mental models of the cardiovascular system (specifically, cardiac output). The study characterized the development of understanding of the system as a function of expertise. The research also documented various conceptual flaws in subjects’ models and how these flaws impacted subjects’ predictions and explanations of physiological manifestations. Figure 4.3 illustrates the four chambers of the heart and blood flow in the pulmonary and cardiovascular systems. The claim is that clinicians and medical students have variably robust representations of the structure and function of the system. This model enables prediction and explanation of the effects of perturbations in the system on blood flow and on various clinical measures such as left ventricular ejection fraction.

*Conceptual* and *procedural* knowledge provide another useful way of distinguishing the functions of different forms of representation. **Conceptual knowledge** refers to one’s understanding of domain-specific concepts. **Procedural knowledge** is a kind of knowing related to how to perform various activities. There are numerous technical skills in medical contexts that necessitate 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 (from reading texts to conversing with colleagues). Procedural knowledge is developed as a function of deliberate practice that results in a learning process known as *knowledge compilation* (Anderson 1983). However, the development of skills may involve a transition from a declarative or interpretive stage



**Fig. 4.3** Schematic model of circulatory and cardiovascular physiology. The diagram illustrates various structures of the pulmonary and systemic circulation system and the process of blood flow. The illustration is used to exemplify the concept of mental model and how it could be applied to explaining and predicting physiologic behavior

toward increasingly proceduralized stages. For example, in learning to use an electronic health record (EHR) system designed to be used as part of a consultation, a less experienced user will need to attend carefully to every action and input, whereas, a more experienced user of this system can more effortlessly interview a patient and simultaneously record patient data (Kushniruk et al. 1996; Patel et al. 2000b). Procedural knowledge supports more efficient and automated action, but is often used without conscious awareness.

Procedural knowledge is often modeled in cognitive science and in artificial intelligence as a production *rule*, which is a condition-action rule that states “if the conditions are satisfied,


then execute the specified action” (either an inference or overt behavior). Production rules are a common method for representing knowledge in medical expert systems such as MYCIN (Davis et al. 1977).

In addition to differentiating between procedural and conceptual knowledge, one can differentiate **factual knowledge** 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 routinely disseminated through a range of sources such as pamphlets and websites. The acquisition of factual knowledge alone is not likely to lead to any increase in understanding or behavioral change (Bransford et al. (1999). 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. This is contrast to a new medical student who may have largely factual knowledge.

Thus far, we have only considered domain-general ways of characterizing the organization of knowledge. In view to understand the nature of medical cognition, it is necessary to characterize the domain-specific nature of knowledge organization in medicine. Given the vastness and complexity of the domain of medicine, this can be a rather daunting task. Clearly, there is no single way to represent all biomedical (or even clinical) knowledge, but it is an issue of considerable importance for research in biomedical informatics. Much research has been conducted in biomedical artificial intelligence with the aim of developing biomedical ontologies for use in knowledge-based systems. Patel et al. (1997) address this issue in the context of using empirical evidence from psychological experiments on medical expertise to test the validity of the AI systems. **Biomedical taxonomies, nomenclatures and vocabulary** systems such as UMLS or SNOMED (see Chap. 7) are engaged in a similar pursuit.

We have employed an epistemological framework developed by Evans and Gadd (1989). They

**Fig. 4.4** Epistemological frameworks representing the structure of medical knowledge for problem solving



proposed a framework that serves to characterize the knowledge used for medical understanding and problem solving, and also for differentiating the levels at which biomedical knowledge may be organized. This framework represents a formalization of biomedical knowledge as realized in textbooks and journals, and can be used to provide us with insight into the organization of clinical practitioners' knowledge (see Fig. 4.4).

The framework consists of a hierarchical structure of concepts formed by *clinical observations* at the lowest level, followed by *findings*, *facets*, and *diagnoses*. Clinical 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 are composed of 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 pathological descriptions such as left-ventricular failure or thyroid condition. Facets resemble the kinds of constructs used by researchers in medical artificial intelligence to describe the partitioning of a problem space. 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 lev-

els beneath it. Finally, the systems level consists of information that serves to contextualize a particular problem, such as the ethnic background of a patient.

#### 4.4 Medical Cognition

The study of expertise is one of the principal paradigms in problem-solving research. Comparing experts to novices provides us with the opportunity to explore the aspects of performance that undergo change and result in increased problem-solving skill (Lesgold 1984; Glaser 2000). It also permits investigators to develop domain-specific models of competence that can be used for assessment and training purposes.

A goal of this approach has been to characterize expert performance in terms of the knowledge and cognitive processes used in comprehension, problem solving, and decision making, using carefully developed laboratory tasks (Chi and Glaser 1981), (Lesgold et al. 1988). deGroot's (1965) pioneering research in chess represents one of the earliest characterizations of expert-novice differences. In one of his experiments, subjects were allowed to view a chess board for 5–10 seconds and were then required to reproduce the position of the chess pieces from memory. The grandmaster chess players were able to reconstruct the mid-game positions with better than 90 % accuracy, while novice chess players could only reproduce approximately 20 % of the correct positions. When the chess pieces were

placed on the board in a random configuration, not encountered in the course of a normal chess match, expert chess masters' recognition ability fell to that of novices. This result suggests that superior recognition ability is not a function of superior memory, but is a result of an enhanced ability to recognize typical situations (Chase and Simon 1973). This phenomenon is accounted for by a process known as "chunking." It is the most general representational construct that makes the fewest assumptions about cognitive processing.


It is well known that knowledge-based differences impact the problem representation and determine the strategies a subject uses to solve a problem. Simon and Simon (1978) compared a novice subject with an expert subject in solving textbook physics problems. The results indicated that the expert solved the problems in one quarter of the time required by the novice with fewer errors. The novice solved most of the problems by working *backward* from the unknown problem solution to the givens of the problem statement. The expert worked *forward* from the givens to solve the necessary equations and determine the particular quantities they are asked to solve for. Differences in the directionality of reasoning by levels of expertise has been demonstrated in diverse domains from computer programming (Perkins et al. 1990) to medical diagnoses (Patel and Groen 1986).

The expertise paradigm spans the range of content domains including physics (Larkin et al. 1980), sports (Allard and Starkes 1991), music (Sloboda 1991), and medicine (Patel et al. 1994). Edited volumes (Ericsson 2006; Chi et al. 1988; Ericsson and Smith 1991; Hoffman 1992) provide an informative general overview of the area. This research has focused on differences between subjects varying in levels of expertise in terms of memory, reasoning strategies, and in particular the role of domain specific knowledge. Among the expert's characteristics uncovered by this research are the following: (1) experts are capable of perceiving large patterns of meaningful information in their domain, which novices cannot perceive; (2) they are fast at processing and at deployment of different skills required for problem solving; (3) they have superior short-term

and long-term memories for materials (e.g., clinical findings in medicine) within their domain of expertise, but not outside of it; (4) they typically represent problems in their domain at deeper, more principled levels whereas novices show a superficial level of representation; (5) they spend more time assessing the problem prior to solving it, while novices tend to spend more time working on the solution itself and little time in problem assessment; (6) individual experts may differ substantially in terms of exhibiting these kinds of performance characteristics (e.g., superior memory for domain materials).

Usually, someone is designated as an expert based on a certain level of performance, as exemplified by Elo ratings in chess; by virtue of being certified by a professional licensing body, as in medicine, law, or engineering; on the basis of academic criteria, such as graduate degrees; or simply based on years of experience or peer evaluation (Hoffman et al. 1995). The concept of an expert, however, refers to an individual who surpasses competency in a domain (Sternberg and Horvath 1999). Although competent performers, for instance, may be able to encode relevant information and generate effective plans of action in a specific domain, they often lack the speed and the flexibility that we see in an expert. A domain expert (e.g., a medical practitioner) possesses an extensive, accessible knowledge base that is organized for use in practice and is tuned to the particular problems at hand. In the study of medical expertise, it has been useful to distinguish different types of expertise.

Patel and Groen (1991a, b) distinguished between general and specific expertise, a distinction supported by research indicating differences between subexperts (i.e., experts physicians who solve a case outside their field of specialization) and experts (i.e., domain specialist) in terms of reasoning strategies and organization of knowledge. General expertise corresponds to expertise that cuts across medical subdisciplines (e.g., general medicine). Specific expertise results from detailed experience within a medical subdomain, such as cardiology or endocrinology. An individual may possess both, or only generic expertise.




**Fig. 4.5** Schematic representation of intermediate effect. 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 concepts recalled, number of conceptual elaborations, or number of hypotheses generated in a variety of tasks

The development of expertise can follow a somewhat unusual trajectory. It is often assumed that the path from novice to expert goes through a steady process of 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, quality) gradually increases. However, research has shown that this assumption is often incorrect (Lesgold et al. 1988; Patel et al. 1994). Cross-sectional studies of experts, intermediates, and novices have shown that people at intermediate levels of expertise may perform more poorly than those at lower level of expertise on some tasks. Furthermore, there is a longstanding body of research on learning that has suggested that the learning process involves phases of error-filled performance followed by periods of stable, relatively error-free performance. In other words, human learning does not consist of the gradually increasing accumulation of knowledge and fine-tuning of skills. Rather, it requires the arduous process of continually learning, re-learning, and exercising new knowledge, punctuated by periods of apparent decrease in mastery and declines in performance, which may be necessary for learning to take place. Figure 4.5 presents an illustration of this learning and developmental phenomenon known as the **intermediate effect**.

The intermediate effect has been found in a variety of tasks and with a great number of performance indicators. The tasks used include comprehension and explanation of clinical problems, doctor-patient communication, recall and explanation of laboratory data, generation of diagnostic hypotheses, and problem solving (Patel and Groen 1991a, b). The performance indicators used have included recall and inference of medical-text information, recall and inference of diagnostic hypotheses, generation of clinical findings from a patient in doctor-patient interaction, and requests for laboratory data, among others. The research has also identified developmental levels at which the intermediate phenomenon occurs, including senior medical students and residents. It is important to note, however, that in some tasks, the development is **monotonic**. For instance, in diagnostic accuracy, there is a gradual increase, with an intermediate exhibiting higher degree of accuracy than the novice and the expert demonstrating a still higher degree than the intermediate. Furthermore, when relevancy of the stimuli to a problem is taken into account, an appreciable monotonic phenomenon appears. For instance, in recall studies, novices, intermediates, and experts are assessed in terms of the total number of propositions recalled showing the typical non-monotonic effect. However, when propositions are divided in terms of their relevance to the problem (e.g., a clinical case), experts recall more relevant propositions than intermediates and novices, suggesting that intermediates have difficulty separating what is relevant from what is not.


During the periods when the intermediate effect occurs, a reorganization of knowledge and skills takes place, characterized by shifts in perspectives or a realignment or creation of goals. The intermediate effect is also partly due to the unintended changes that take place as the person reorganizes for intended changes. People at intermediate levels typically generate a great deal of irrelevant information and seem incapable of discriminating what is relevant from what is not. As compared to a novice student (Fig. 4.6), the reasoning pattern of an intermediate student shows the generation of long chains of discussion



**Fig. 4.6** Problem interpretations by a novice medical student. The given information from patient problem is represented on the *right side of the figure* and the new generated information is given on the *left side*, information in the box represents diagnostic hypothesis. Intermediate hypothesis are represented as *solid dark*

*circles (filled)*. Forward driven or data driven inference *arrows* are shown from left to right (*solid dark line*). Backward or hypothesis driven inference *arrows* are shown from right to left (*solid light line*). *Thick solid dark line* represents rule out strategy

**Fig. 4.7** Problem interpretations by an intermediate medical student




evaluating multiple hypotheses and reasoning in haphazard direction (Fig. 4.7). A well-structured knowledge structure of a senior level student leads him more directly to a solution (Fig. 4.8). Thus, the intermediate effect can be explained as a function of the learning process, maybe as a necessary phase of learning. Identifying the factors involved in the intermediate effect may help in improving performance during learning (e.g., by designing decision-support systems or

intelligent tutoring systems that help the user in focusing on relevant information).

There are situations, however, in which the intermediate effect disappears. Schmidt reported that the intermediate recall phenomenon disappears when short text-reading times are used. Novices, intermediates, and experts given only a short time to read a clinical case (about thirty seconds) recalled the case with increasing accuracy. This suggests that under time-restricted



**Fig. 4.8** Problem interpretations by a senior medical student



conditions, intermediates cannot engage in extraneous search. In other words, intermediates that are not under time pressure process too much irrelevant information whereas experts do not. On the other hand, novices lack the knowledge to do much searching. Although intermediates may have most of the pieces of knowledge in place, this knowledge is not sufficiently well organized to be efficiently used. Until this knowledge becomes further organized, the intermediate is more likely to engage in unnecessary search.

The intermediate effect is not a one-time phenomenon. Rather, it occurs repeatedly at strategic points in a student or physician's training and follow periods in which large bodies of new knowledge or complex skills are acquired. These periods are followed by intervals in which there is a decrement in performance until a new level of mastery is achieved.

#### 4.4.1 Expertise in Medicine

The systematic investigation of medical expertise began more than 50 years ago with research by Ledley and Lusted (1959) 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. The earliest empirical studies of medical expertise can be traced to the works of Rimoldi (1961) and Kleinmuntz (1968) who conducted experimental studies of diagnostic reasoning by 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).

The origin of contemporary research on medical thinking is associated with the seminal work of Elstein, Shulman, and Sprafka (1978) who studied the problem solving processes of physicians by drawing on then contemporary methods and theories of cognition. This model of problem solving has had a substantial influence both on studies of medical cognition 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, physicians were selective in the data they collected, focusing only on the relevant data. Third, physicians made use of a hypothetico-deductive method of diagnostic reasoning. The hypothetico-deductive process was viewed as consisting of four stages: cue acquisition, hypothesis generation, cue interpretation, and hypothesis evaluation. Attention to initial cues led to the rapid generation of a few select hypotheses. According to the authors, each cue was interpreted as positive, negative or non-contributory to each hypothesis generated. They were unable to find differences their diagnostic reasoning strategies between superior physicians (as judged by their peers) and other physicians (Elstein et al. 1978).


The previous research was largely modeled after early problem-solving studies in knowledge-lean tasks. Medicine is clearly a knowledge-rich domain and a different approach was needed. Feltovich, Johnson, Moller, and Swanson (1984), 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. For example, novice's knowledge was described as "classically-centered", built around the prototypical instances of a disease category. The disease models were described as sparse and lacking cross-referencing between shared features of disease categories in memory. In contrast, experts' memory store of disease models was found to be extensively cross-referenced with a rich network of connections among diseases that can present with similar symptoms. These differences accounted for subjects' inferences about diagnostic cues and evaluation of competing hypotheses.

Patel and colleagues studied the knowledge-based solution strategies of expert cardiologists as evidenced by their pathophysiological explanations of a complex clinical problem (Patel and Groen 1986). 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 is in contrast to subjects who misdiagnosed or partially diagnosed the patient problem. They 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, Shulman, and Sprafka (1978), which did not differentiate expert from non-expert reasoning strategies.

Patel and Groen (1991a, b) investigated the nature and directionality of clinical reasoning in a range of contexts of varying complexity. The objectives of this research program were both to advance our understanding of medical expertise and to devise more effective ways of teaching clinical problem solving. It has been established that 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 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 reasoning have been found in other domains (Larkin et al. 1980). Due to 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. 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



**Fig. 4.9** Diagrammatic representation of data-driven (*top down*) and hypothesis-driven (*bottom-up*) reasoning. From the presence of vitiligo, a prior history of progressive thyroid disease, and examination of the thyroid (clinical findings on the *left side of figure*), the physician reasons forward to conclude the diagnosis of Myxedema

(*right of figure*). 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* refers to a conditional relation, *CAU* indicates a causal relation, and *RSLT* identifies a resultive relation

the initial problem statement to the problem solution, as illustrated in Fig. 4.9. In contrast, hypothesis-driven reasoning is slower and may make heavy demands on working memory, 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 usually exemplary of a *weak method* of problem solving in the sense that 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 have shown that the pattern of data-driven reasoning breaks down in conditions of case complexity, unfamiliarity with the problem, and uncertainty (Patel et al. 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.

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 training in the use of problem solving strategies, the situation produces a heavy load on cognitive resources and 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 to solve the problem-solving method and learning the content of the material. It has been found that when subjects used a strategy based on the use of data-driven reasoning, they were 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 (Patel et al. 1990).

Visual diagnosis has also been an active area of inquiry in medical cognition. Studies have investigated clinicians at varying levels of expertise in their ability to diagnose skin lesions presented on a slide. 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, that reflected the underlying pathophysiological structure.

The ability to abstract the underlying principles of a problem is considered to be one of the hallmarks of expertise, both in medical problem solving and in other domains (Chi and Glaser 1981). Lesgold et al. (1988) investigated the abilities of radiologists at different levels of expertise, in the interpretation of 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 anatomical localization and served to constrain the possible interpretations. Novices experienced greater difficulty focusing in on the important structures and were more likely to maintain inappropriate interpretations despite discrepant findings in the patient history.

Crowley, Naus, Stewart, and Friedman (2003) employed a similar protocol-analytic approach to the Lesgold study to examine differences in expertise in breast pathology. The results suggest systematic differences between subjects at varying levels of expertise corresponding to accuracy of diagnosis, and all aspects of task performance including microscopic search, feature detection, feature identification and data interpretation. The authors propose a model of visual diagnostic competence that involves development of effective search strategies, fast and accurate recognition of anatomic location, acquisition of visual data interpretation skills and explicit feature identification strategies that results from a well-organized knowledge base.

The study of medical cognition has been summarized in a series of articles (Patel et al. 1994) and edited volumes (e.g., Evans and Patel 1989). Other active areas of research include medical text comprehension, therapeutic reasoning and mental models of physiological systems. Medical cognition remains an active area of research and continues to inform debates regarding medical

curricula and approaches to learning (Patel et al. 2005; Schmidt and Rikers 2007).

---

## 4.5 Human Computer Interaction

**Human computer interaction (HCI)** is a multifaceted discipline devoted to the study and practice of design and usability (Carroll 2003). The history of computing and more generally, the history of artifacts design, are rife with stories of dazzlingly powerful devices with remarkable capabilities that are thoroughly unusable by anyone except for the team of designers and their immediate families. In the *Psychology of Everyday Things*, Norman (1988) describes a litany of poorly designed artifacts ranging from programmable VCRs to answering machines and water faucets that are inherently non-intuitive and very difficult to use. Similarly, there have been numerous innovative and promising medical information technologies that have yielded decidedly sub-optimal results and deep user dissatisfaction when implemented in practice. At minimum, difficult interfaces result in steep learning curves and structural inefficiencies in task performance. At worst, problematic interfaces can have serious consequences for patient safety (Lin et al. 1998; Zhang et al. 2004; Koppel et al. 2005) (see Chap. 11).

Twenty years ago, Nielsen (1993) reported that around 50 % of software code was devoted to the user interface and a survey of developers indicated that, on average, 6 % of their project budgets were spent on usability evaluation. Given the ubiquitous presence of graphical user interfaces (GUI), it is likely that more than 50 % of code is now devoted to the GUI. On the other hand, usability evaluations have greatly increased over the course of the last 10 years (Jaspers 2009). There have been numerous texts devoted to promoting effective user interface design (Preece et al. 2007; Shneiderman 1998) and the importance of enhancing the user experience has been widely acknowledged by both consumers and producers of information technology (see Chap. 11). Part of the impetus is

that usability has been demonstrated to be highly cost effective. Karat (1994) reported that for every dollar a company invests in the usability of a product, it receives between \$10 and \$100 in benefits. Although much has changed in the world of computing since Karat's estimate (e.g., the flourishing of the World Wide Web), it is very clear that investments in usability still yield substantial rates of return (Nielsen et al. 2008). It remains far more costly to fix a problem after product release than in an early design phase. In our view, usability evaluation of medical information technologies has grown substantially in prominence. The concept of usability as well as the methods and tools to measure and promote it are now "touchstones in the culture of computing" (Carroll 2003).

Usability methods have been used to evaluate a wide range of medical information technologies including infusion pumps (Dansky et al. 2001), ventilator management systems, physician order entry (Ash et al. 2003a; Koppel et al. 2005), pulmonary graph displays (Wachter et al. 2003), information retrieval systems, and research web environments for clinicians (Elkin et al. 2002). In addition, usability techniques are increasingly used to assess patient-centered environments (Cimino et al. 2000; Kaufman et al. 2003; Chan and Kaufman 2011). The methods include observations, focus groups, surveys and experiments. Collectively, these studies make a compelling case for the instrumental value of such research to improve efficiency, user acceptance and relatively seamless integration with current workflow and practices.

What do we mean by usability? Nielsen (1993) suggests that usability includes the following five attributes: (1) *learnability*: system should be relatively easy to learn, (2) *efficiency*: an experienced user can attain a high level of productivity, (3) *memorability*: features supported by the system should be easy to retain once learned, (4) *errors*: system should be designed to minimize errors and support error detection and recovery, and (5) *satisfaction*: the user experience should be subjectively satisfying.

Even with growth of usability research, there remain formidable challenges to designing and developing usable systems. This is exemplified

by the events at Cedar Sinai Medical Center in which a decision was made to suspend use of a computer-based physician order entry system just a few months after implementation. Physicians complained that the system, which was designed to reduce medical errors, compromised patient safety, took too much time and was difficult to use (Benko 2003). To provide another example, we have been working with a mental health computer-based patient record system that is rather comprehensive and supports a wide range of functions and user populations (e.g., physicians, nurses, and administrative staff). However, clinicians find it exceptionally difficult and time-consuming to use. The interface is based on a form (or template) metaphor and is neither user- nor task-centered. The interface emphasizes completeness of data entry for administrative purposes rather than the facilitation of clinical communication and is not optimally designed to support patient care (e.g., efficient information retrieval and useful summary reports). In general, the capabilities of this system are not readily usefully deployed to improve human performance. We further discuss issues of EHRs in a subsequent section.


Innovations in technology guarantee that usability and interface design will be a perpetually moving target. In addition, as health information technology reaches out to populations across the digital divide (e.g., seniors and low literacy patient populations), there is a need to consider new interface requirements. Although evaluation methodologies and guidelines for design yield significant contributions, there is a need for a scientific framework to understand the nature of user interactions. **Human computer interaction (HCI)** is a multifaceted discipline devoted to the study and practice of usability. HCI has emerged as a central area of both computer science, development and applied social science research (Carroll 2003).

HCI has spawned a professional orientation that focuses on practical matters concerning the integration and evaluation of applications of technology to support human activities. There are also active academic HCI communities that have contributed significant advances to the science of computing. HCI researchers have been devoted

to the development of innovative design concepts such as **virtual reality**, **ubiquitous computing**, **multimodal interfaces**, **collaborative workspaces**, and **immersive environments**. HCI research has been instrumental in transforming the software engineering process towards a more user-centered iterative system development (e.g., rapid prototyping). HCI research has also been focally concerned with the cognitive, social, and cultural dimensions of the computing experience. In this regard, it is concerned with developing analytic frameworks for characterizing how technologies can be used more productively across a range of tasks, settings, and user populations.

Carroll (1997) traces the history of HCI back to the 1970s with the advent of **software psychology**, a behavioral approach to understanding and furthering software design. Human factors, ergonomics, and industrial engineering research were pursuing some of the same goals along parallel tracks. In the early 1980s, Card et al. (1983) envisioned HCI as a test bed for applying cognitive science research and also furthering theoretical development in cognitive science. The Goals, Operators, Methods, and Selection Rules (GOMS) approach to modeling was a direct outgrowth of this initiative. GOMS is a powerful predictive tool, but it is limited in scope to the analysis of routine skills and expert performance. Most medical information technologies such as provider-order entry systems, engage complex cognitive skills.

HCI research has embraced a diversity of approaches with an abundance of new theoretical frameworks, design concepts, and analytical foci (Rogers 2004). Although we view this as an exciting development, it has also contributed to a certain scientific fragmentation (Carroll 2003). Our own research is grounded in a cognitive engineering framework, which is an interdisciplinary approach to the development of principles, methods and tools to assess and guide the design of computerized systems to support human performance (Roth et al. 2002). In supporting performance, the focus is on cognitive functions such as attention, perception, memory, comprehension, problem solving, and decision making. The approach is centrally concerned with the analysis of cognitive tasks and



**Fig. 4.10** Norman's seven stage model of action

the processing constraints imposed by the human cognitive system.

Models of **cognitive engineering** are typically predicated on a cyclical pattern of interaction with a system. This pattern is embodied in Norman's (1986) seven stage model of action, illustrated in Fig. 4.10. The action cycle begins with a *goal*, for example, retrieving a patient's medical record. The goal is abstract and independent of any system. In this context, let's presuppose that the clinician has access to both a paper record and an electronic record. The second stage involves the formation of an *intention*, which in this case might be to retrieve the record online. The intention leads to the *specification of an action* sequence, which may include logging onto the system (which in itself may necessitate several actions), engaging a search facility to retrieve information, and entering the patient's medical record number or some other identifying information. The specification results in *executing an action*, which may necessitate several behaviors. The system responds in some fashion (or doesn't respond at all). The user may or may not perceive a change in system state (e.g., system provides no indicators of a wait state). The

perceived system response must then be *interpreted* and *evaluated* to determine whether the goal has been achieved. This will then determine whether the user has been successful or whether an alternative course of action is necessary.

A complex task will involve substantial nesting of subgoals, involving a series of actions that are necessary before the primary goal can be achieved. To an experienced user, the action cycle may appear to be completely seamless. However to a novice user, the process may break down at any of the seven stages. There are two primary means in which the action cycle can break down. The *gulf of execution* reflects the difference between the goals and intentions of the user and the kinds of actions enabled by the system. A user may not know the appropriate action sequence or the interface may not provide the prerequisite features to make such sequences transparent. For example, many systems require a goal completion action, such as pressing “Enter”, after the primary selection had been made. This is a source of confusion, especially for novice users. The *gulf of evaluation* reflects the degree to which the user can interpret the state of the system and determine how well their expectations have been met. For example, it is sometimes difficult to interpret a state transition and determine whether one has arrived at the right place. Goals that necessitate multiple state or screen transitions are more likely to present difficulties for users, especially as they learn the system. Bridging gulfs involves both bringing about changes to the system design and educating users to foster competencies that can be used to make better use of system resources.

Gulfs are partially attributable to differences in the designer’s model and the users’ mental models. The designer’s model is the conceptual model of intent of the system, partially based on an estimation of the user population and task requirements (Norman 1986). The users’ mental models of system behavior are developed through interacting with similar systems and gaining an understanding of how actions (e.g., clicking on a link) will produce predictable and desired outcomes. Graphical user interfaces that involve direct manipulation of screen objects represent an attempt to reduce the distance

between a designer’s and users’ model. The distance is more difficult to close in a system of greater complexity that incorporates a wide range of functions, like many medical information technologies.

Norman’s theory of action has given rise to (or in some case reinforced) the need for sound design principles. For example, the state of a system should be plainly *visible* to the user. There is a need to provide good mappings between the actions (e.g., clicking on a button) and the results of the actions as reflected in the state of the system (e.g., screen transitions). Similarly, a well-designed system will provide full and continuous feedback so that the user can understand whether one’s expectations have been met.

The model has also informed a range of cognitive task-analytic usability evaluation methods such as the cognitive walkthrough (Polson et al. 1992), described below. 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 performing 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 guidelines to the investigation of EHR systems. Generic tasks necessitate similar cognitive demands and have a common underlying structure that involves similar kinds of reasoning and patterns of inference. For example, clinical tasks in medicine include diagnostic reasoning, therapeutic reasoning, and patient monitoring and management. Similarly, an admission order-entry task can be completed using written orders or one of many diverse computer-based order-entry systems. The underlying task of communicating orders in view to admit a patient remains the same. However, the particular implementation will greatly impact the performance of the task. For example, a system may eliminate the need for redundant entries and greatly facilitate the process. On the other hand, it may introduce unnecessary complexity leading to suboptimal performance.

These are a class of usability evaluation methods performed by expert analysts or reviewers and unlike usability testing, don’t typically

involve the use of subjects (Nielsen 1994)). The **cognitive walkthrough** (CW) and the heuristic evaluation are the most commonly used inspection methods. **Heuristic evaluation** (HE) is a method by which an application is evaluated on the basis of a small set of well-tested design principles such as visibility of system status, user control and freedom, consistency and standards, flexibility and efficiency of use (Nielsen 1993). We illustrate HE in the context of a human factors study later in the chapter. The CW is a cognitive task analytic method that has been applied to the study of usability and learnability of several distinct medical information technologies (Kushniruk et al. 1996). The purpose of a CW is to characterize the cognitive processes of users performing a task. The method involves identifying sequences of actions and goals needed to accomplish a given task. The specific aims of the procedure are to determine whether the typical user's background knowledge and the cues generated by the interface are likely to be sufficient to produce the correct goal-action sequence required to perform a task. The method is intended to identify potential usability problems that may impede the successful completion of a task or introduce complexity in a way that may frustrate real users. The method is performed by an analyst or group of analysts 'walking through' the sequence of actions necessary to achieve a goal. Both behavioral or physical actions such as mouse clicks and cognitive actions (e.g., inference needed to carry out a physical action) are coded. The principal assumption underlying this method is that a given task has a specifiable goal-action structure (i.e., the ways in which a user's objectives can be translated into specific actions). As in Norman's model, each action results in a system response (or absence of one), which is duly noted.

The CW method assumes a cyclical pattern of interaction as described previously. The codes for analysis include *goals*, which can be decomposed into a series of *subgoals* and *actions*. For example, opening an Excel spreadsheet (goal) may involve locating an icon or shortcut on one's desktop (subgoal) and double clicking on the application (action). The system response (e.g.,

change in screen, update of values) is also characterized and an attempt is made to discern potential problems. This is illustrated below in a partial walkthrough of an individual obtaining money from an automated teller system.

**Goal:** Obtain \$80 Cash from Checking Account

1. **Action:** Enter Card (Screen 1)

**System response:** Enter PIN>(Screen 2)

2. **Subgoal:** Interpret prompt and provide input

3. **Action:** Enter "Pin" on Numeric keypad

4. **Action:** Hit Enter (press lower white button next to screen)

**System response:** "Do you Want a Printed Transaction Record".

Binary Option: Yes or No (Screen 3)

5. **Subgoal:** Decide whether a printed record is necessary

6. **Action:** Press Button Next to No Response

**System response:** Select Transaction-8 Choices (Screen 4)

7. **Subgoal:** Choose Between Quick Cash and Cash Withdrawal

8. **Action:** Press Button Next to Cash Withdrawal

**System response:** Select Account (Screen 5)

9. **Action:** Press Button Next to Checking

**System response:** Enter Dollar Amounts in Multiples of 20 (Screen 6)

10. **Action:** Enter \$80 on Numeric Key Pad

11. **Action:** Select Correct

The walkthrough of the ATM reveals that process of obtaining money from the ATM necessitated a minimum of eight actions, five goals and subgoals, and six screen transitions. In general, it is desirable to minimize the number of actions necessary to complete a task. In addition, multiple screen transitions are more likely to confuse the user. We have employed a similar approach to analyze the complexity of a range of medical information technologies including EHRs, a home telecare system, and infusion pumps used in intensive care settings. The CW process emphasizes the sequential process, not unlike problem solving, involved in completing a computer-based task. The focus is more on the process rather than the content of the displays.



**Usability testing** represents the gold standard in usability evaluation methods. It refers to a class of methods for collecting empirical data of representative users performing representative tasks. It is known to capture a higher percentage of the more serious usability problems and provides a greater depth of understanding into the nature of the interaction (Jaspers 2009). Usability testing commonly employs video capture of users performing the tasks as well as video-analytic methods of analysis (Kaufman et al. 2003; Kaufman et al. 2009). It involves in-depth testing of a small number of subjects. The assumption is that a test can be perfectly valid with as few as five or six subjects. In addition, five or six subjects may find upwards of 80 % of the usability problems. A typical usability testing study will involve five to ten subjects who are asked to think aloud as they perform the task.

It's not uncommon to employ multiple methods such as the cognitive walkthrough and usability testing (Beuscart-Zephir et al. 2005a, b). The methods are complementary and serve as a means to triangulate significant findings. Kaufman et al. (2003) conducted a cognitive evaluation of the IDEATel home telemedicine system (Shea et al. 2002; 2009); Starren et al. 2002; Weinstock et al. 2010) with a particular focus on a) system usability and learnability, and b) the core competencies, skills and knowledge necessary to productively use the system. The study employed both a cognitive walkthrough and in-depth usability testing. The focal point of the intervention was the home telemedicine unit (HTU), which provided the following functions: (1) synchronous video-conferencing with a nurse case manager, (2) electronic transmission of fingerstick glucose and blood pressure readings, (3) email to a physician and nurse case manager, (4) review of one's clinical data and (5) access to Web-based educational materials (see Chap. 18 for more details on IDEATel). The usability study revealed dimensions of the interface that impeded optimal access to system resources. In addition, significant obstacles corresponding to perceptual-motor skills, mental models of the system, and health literacy were documented.

## 4.6 Human Factors Research and Patient Safety

Human error in medicine, and the adverse events which may follow, are problems of psychology and engineering not of medicine "(Senders, 1993)" (cited in (Woods et al. 2007).

**Human factors** research is a discipline devoted to the study of technology systems and how people work with them or are impacted by these technologies (Henriksen 2010). Human factors research discovers and applies information about human behavior, abilities, limitations, and other characteristics to the design of tools, machines, systems, tasks, and jobs, and environments for productive, safe, comfortable, and effective human use (Chapanis 1996). In the context of healthcare, human factors is concerned with the full complement of technologies and systems used by a diverse range of individuals including clinicians, hospital administrators, health consumers and patients (Flin and Patey 2009). Human factors work approaches the study of health practices from several perspectives or levels of analysis. The focus is on the ways in which organizational, cultural, and policy issues inform and shape healthcare processes. A full exposition of human factors in medicine is beyond the scope of this chapter. For a detailed treatment of these issues, the reader is referred to the Handbook of Human Factors and Ergonomics in Health Care and Patient Safety (Carayon 2007). The focus in this chapter is on cognitive work in human factors and healthcare, particularly in relation to issues having to do with patient safety. We recognize that patient safety is a systemic challenge at a multiple levels of aggregation beyond the individual. It is clear that understanding, predicting and transforming human performance in any complex setting requires a detailed understanding of both the setting and the factors that influence performance (Woods et al. 2007).

Our objective in this section is to introduce a theoretical foundation, establish important concepts and discuss illustrative research in **patient safety**. Human factors and human computer

interaction are different disciplines with different histories and different professional and academic societies. HCI is more focused on computing and cutting-edge design and technology, whereas human factors focus on a broad range of systems that include, but are not restricted to computing technologies (Carayon 2007). Patient safety is one of the central issues in human factors research and we address this issue in greater detail in a subsequent section. Both human factors and HCI employ many of the same methods of evaluation and both strongly emphasize a user-centered approach to design and a systems-centered approach to the study of technology use. Researchers and professionals in both domains draw on certain theories including cognitive engineering. The categorization of information technology-based work as either human factors or HCI is sometimes capricious.

The field of human factors is guided by principles of engineering and applied cognitive psychology (Chapanis 1996). Human factors analysis applies knowledge about the strengths and limitations of humans to the design of interactive systems, equipment, and their environment. The objective is to ensure their effectiveness, safety, and ease of use. Mental models and issues of decision making are central to human factors analysis. Any system will be easier and less burdensome to use to the extent that that it is co-extensive with users' mental models. Human factors focus on different dimensions of cognitive capacity, including memory, attention, and workload. Our perceptual system inundates us with more stimuli than our cognitive systems can possibly process. Attentional mechanisms enable us to selectively prioritize and attend to certain stimuli and attenuate other ones. They also have the property of being sharable, which enables us to multitask by dividing our attention between two activities. For example, if we are driving on a highway, we can easily have a conversation with a passenger at the same time. However, as the skies get dark or the weather changes or suddenly you find yourself driving through winding mountainous roads, you will have to allocate more of your attentional resources to driving and less to the conversation.

Human factors research leverages theories and methods from cognitive engineering to characterize human performance in complex settings and challenging situations in aviation, industrial process control, military command control and space operations (Woods et al. 2007). The research has elucidated empirical regularities and provides explanatory concepts and models of human performance. This enables us to discern common underlying patterns in seemingly disparate settings (Woods et al. 2007).

### 4.6.1 Patient Safety

When human error is viewed as a cause rather than a consequence, it serves as a cloak for our ignorance. By serving as an end point rather than a starting point, it retards further understanding. (Henriksen 2008)

Patient safety refers to the prevention of health-care errors, and the elimination or mitigation of patient injury caused by healthcare errors (Patel and Zhang 2007). It has been an issue of considerable concern for the past quarter century, but the greater community was galvanized by the Institute of Medicine Report, "To Err is Human," released in 1999. This report communicated the surprising fact that 98,000 preventable deaths every single year in the United States are attributable to human error, which makes it the 8th leading cause of death in this country. Although one may argue over the specific numbers, there is no disputing that too many patients are harmed or die every year as a result of human actions or absence of action.

The Harvard Medical Practice Study was published several years prior to the IOM report and was a landmark study at the time. Based on an extensive review of patient charts in New York State, they were able to determine that an adverse event occurred in almost 4 % of the cases (Leape et al. 1991). An adverse event refers to any unfavorable change in health or side effect that occurs in a patient who is receiving the treatment. They further determined that almost 70 % of these adverse events were caused by errors and 25 % of all errors were due to negligence.

We can only analyze errors after they happened and they often seem to be glaring blunders after the fact. This leads to assignment of blame or search for a single cause of the error. However, in hindsight, it is exceedingly difficult to recreate the situational context, stress, shifting attention demands and competing goals that characterized a situation prior to the occurrence of an error. This sort of retrospective analysis is subject to hindsight bias. **Hindsight bias** masks the dilemmas, uncertainties, demands and other latent conditions that were operative prior to the mishap. Too often the term ‘human error’ connotes blame and a search for the guilty culprits, suggesting some sort of human deficiency or irresponsible behavior. Human factors researchers recognized that this approach error is inherently incomplete and potentially misleading. They argue for the need for a more comprehensive systems-centered approach that recognizes that error could be attributed to a multitude of factors as well as the interaction of these factors. Error is the failure of a planned sequence of mental or physical activities to achieve its intended outcome when these failures cannot be attributed to chance (Arocha et al. 2005; Reason 1990). Reason (1990) introduced an important distinction between **latent** and **active failures**. Active failure represents the face of error. The effects of active failure are immediately felt. In healthcare, active errors are committed by providers such as nurses, physicians, or pharmacists who are actively responding to patient needs at the “sharp end”. The latent conditions are less visible, but equally important. Latent conditions are enduring systemic problems that may not be evident for some time, combine with other system problems to weaken the system’s defenses and make errors possible. There is a lengthy list of potential latent conditions including poor interface design of important technologies, communication breakdown between key actors, gaps in supervision, inadequate training, and absence of a safety culture in the workplace—a culture that emphasizes safe practices and the reporting of any conditions that are potentially dangerous.

Zhang, Patel, Johnson, and Shortliffe (2004) have developed a taxonomy of errors partially

based on the distinctions proposed by Reason (1990). We can further classify errors in terms of **slips** and **mistakes** (Reason 1990). A slip occurs when the actor selected the appropriate course of action, but it was executed inappropriately. A mistake involves an inappropriate course of action reflecting an erroneous judgment or inference (e.g., a wrong diagnosis or misreading of an x-ray). Mistakes may either be knowledge-based owing to factors such as incorrect knowledge or they may be rule-based, in which case the correct knowledge was available, but there was a problem in applying the rules or guidelines. They further characterize medical errors as a progression of events. There is a period of time when everything is operating smoothly. Then an unsafe practice unfolds resulting in a kind of error, but not necessarily leading to an adverse event. For example, if there is a system of checks and balances that is part of routine practice or if there is systematic supervisory process in place, the vast majority of errors will be trapped and defused in this middle zone. If these measures or practices are not in place, an error can propagate and cross the boundary to become an adverse event. At this point, the patient has been harmed. In addition, if an individual is subject to a heavy workload or intense time pressure, then that will increase the potential for an error, resulting in an adverse event.

The notion that human error should not be tolerated is prevalent in both the public and personal perception of the performance of most clinicians. However, researchers in other safety-critical domains have long since abandoned the quest for zero defect, citing it as an impractical goal, and choosing to focus instead on the development of strategies to enhance the ability to recover from error (Morel et al. 2008). Patel and her colleagues conducted empirical investigations into error detection and recovery by experts (attending physicians) and non-experts (resident trainees) in the critical care domain, using both laboratory-based and naturalistic approaches (Patel et al. 2011). These studies show that expertise is more closely tied to ability to detect and recover from errors and not so much to the ability not to make errors. The study results show that both the experts and

non-experts are prone to commit and recover from errors, but experts' ability to detect and recover from knowledge based errors is better than that of trainees. Error detection and correction in complex real-time critical care situations appears to induce certain urgency for quick action in a high alert condition, resulting in rapid detection and correction. Studies on expertise and understanding of the limits and failures of human decision-making are important if we are to build robust decision-support systems to manage the boundaries of risk of error in decision making (Patel and Cohen 2008).

There has been a wealth of studies regarding patient safety and medical errors in a range of contexts. Holden and Karsh (2007) argue that much of the work is atheoretical in nature and that this diminishes the potential generalizability of the lessons learned. They propose a multifaceted theoretical framework incorporating theories from different spheres of research including motivation, decision-making, and social-cognition. They also draw on a sociotechnical approach, which is a perspective that interweaves technology, people, and the social context of interaction for the design of systems. The end result is a model that can be applied to health information technology usage behavior and that guides a set of principles for design and implementation. The authors propose that through iterative testing of the model, the efforts of researchers and practitioners will yield greater success in the understanding, design, and implementation of health information technology.

#### **4.6.2 Unintended Consequences**

It is widely believed that health information technologies have the potential to transform healthcare in a multitude of ways including the reduction of errors. However, it is increasingly apparent that technology-induced errors are deeply consequential and have had deleterious consequences for patient safety. Ash, Stavri, and Kuperman (2003b) were among the first to give voice to this problem in the informatics community. They also endeavored to describe and

enumerate the primary kinds of errors caused by health information systems, those related to entering and retrieving information and those related to communication and coordination. The authors characterize several problems that are not typically found in usability studies. For example, many interfaces are not suitable for settings that are highly interruptive (e.g., a cluttered display with too many options). They also characterize a problem in which an information entry screen that is highly structured and requires completeness of entry can cause cognitive overload.

Medical devices include any healthcare product, excluding drugs, that are used for the purpose of prevention, diagnosis, monitoring, treatment or alleviation of an illness (Ward and Clarkson 2007). There is considerable evidence that suggests that medical devices can also cause substantial harm (Jha et al. 2010). It has been reported that more than one million adverse medical device events occur annually in the United States (Bright and Brown 2007). Although medical devices are an integral part of medical care in hospital settings, they are complex in nature and clinicians often do not receive adequate training (Woods et al. 2007). In addition, many medical devices such as smart infusion pumps, patient controlled analgesia (PCA) devices, and bar coded medication administration systems have been partially automated and offer a complex programmable interface (Beuscart-Zephir et al. 2005a, 2007). Although this affords opportunities to facilitate clinical care and medical decision making, it may add layers of complexity and uncertainty.

There is evidence to suggest that a poorly designed user interface can present substantial challenges even for the well-trained and highly skilled user (Zhang et al. 2003). Lin and colleagues (1998) conducted a series of studies on a patient controlled analgesic or PCA device, a method of pain relief that uses disposable or electronic infusion devices and allows patients to self-administer analgesic drugs as required. The device is programmed by a nurse or technician and this limits the maximum level of drug administration to keep the dose within safe lev-

els. Lin and colleagues investigated the effects of two interfaces to a commonly used PCA device including the original interface. Based on a **cognitive task analysis**, they redesigned the original interface so that it was more in line with sound human factors principles. As described previously, cognitive task analysis is a method that breaks a task into sets of subtasks or steps (e.g., a single action), the system responses (e.g., changes in the display as a result of an action) and inferences that are needed to interpret the state of the system. It is an effective gauge of the complexity of a system. For example, a simple task that necessitates 25 steps or more to complete using a given system is likely to be unnecessarily complex. On the basis of the cognitive task analysis, they found the existing PCA interface to be problematic in several different ways. For example, the structure of many subtasks in the programming sequence was unnecessarily complex. There was a lack of information available on the screen to provide meaningful feedback and to structure the user experience (e.g., negotiating the next steps). For example, a nurse would not know that he or she was on the third of five screens or when they were half way through the task.

On the basis of the CTA analysis, Lin and colleagues (1998) also redesigned the interface according to sound human factors principles. The new system was designed to simplify the entire process and provide more consistent feedback. It's important to note that the revised screen was a computer simulation and was not actually implemented in the physical device. They conducted a cognitive study with 12 nurses comparing simulations of the old and new interface. They found that programming the new interface was 15 % faster. The average workload rating for the old interface was twice as high. The new interface led to 10 errors as compared to 20 for the old one. This is a compelling demonstration that medical equipment can be made safer and more efficient by adopting sound human factors design principles.

This methodology embodies a particular philosophy that emphasizes simplicity and functionality over intricacy of design and presentation.

Zhang and colleagues employed a modified heuristic evaluation method (see section 4.5, above) to test the safety of two infusion pumps (Zhang et al. 2003). On the basis of an analysis by 4 evaluators, a total of 192 violations with the user interface design were documented. Consistency and visibility (the ease in which a user can discern the system state) were the most widely documented violations. Several of the violations were classified as problems of substantial severity. Their results suggested that one of the two pumps was likely to induce more medical errors than the other ones.

It is clear that usability problems are consequential and have the potential to impact patient safety. Kushniruk et al. (2005) examined the relationship between particular kinds of usability problems and errors in a handheld prescription writing application. They found that particular usability problems were associated with the occurrence of error in entering medication. For example, the problem of inappropriate default values automatically populating the screen was found to be correlated with errors in entering wrong dosages of medications. In addition, certain types of errors were associated with mistakes (not detected by users) while others were associated with slips pertaining to unintentional errors. Horsky et al. (2005) analyzed a problematic medication order placed using a CPOE system that resulted in an overdose of potassium chloride being administered to an actual patient. The authors used a range of investigative methods including inspection of system logs, semi-structured interviews, examination of the electronic health record, and cognitive evaluation of the order entry system involved. They found that the error was due to a confluence of factors including problems associated with the display, the labeling of functions and ambiguous dating of the dates in which a medication was administered. The poor interface design did not provide assistance with the decision-making process, and in fact, its design served as a hindrance, where the interface was a poor fit for the *conceptual operators* utilized by clinicians when calculating medication dosage (i.e., based on volume not duration).

Koppel and colleagues (2005) published an influential study examining the ways in which computer-provider order-entry systems (CPOE) facilitated medical errors. The study, which was published in JAMA (Journal of the American Medical Association), used a series of methods including interviews with clinicians, observations and a survey to document the range of errors. According to the authors, the system facilitated 22 types of medication error and many of them occurred with some frequency. The errors were classified into two broad categories: (1) information errors generated by fragmentation of data and failure to integrate the hospital's information systems and (2) human-machine interface flaws reflecting machine rules that do not correspond to work organization or usual behaviors.

It is a well-known phenomenon that users come to rely on technology and often treat it as an authoritative source that can be implicitly trusted. This can result in information/fragmentation errors. In this case, clinicians relied on CPOE displays to determine the minimum effective dose or a routine dose for a particular kind of patient. However, there was a discrepancy between their expectations and the dose listing. The dosages listed on the display were based on the pharmacy's warehousing and not on clinical guidelines. For example, although normal dosages are 20 or 30 mg, the pharmacy might stock only 10-mg doses, so 10-mg units are displayed on the CPOE screen. Clinicians mistakenly assumed that this was the minimal dose. Medication discontinuation failures are a commonly documented problem with CPOE systems. The system expects a clinician to (1) order new medications and (2) cancel existing orders that are no longer operative. Frequently, clinicians fail to cancel the existing orders leading to duplicative medication orders and thereby increasing the possibility of medical errors. Perhaps, a reminder that prior orders exist and may need to be canceled may serve to mitigate this problem.

As is the case with other clinical information systems, CPOE systems suffer from a range of usability problems. The study describes three kinds of problems. When selecting a patient, it is relatively easy to select the wrong patient because

names and drugs are close together, the font is small, and, patients' names do not appear on all screens. On a similar note, physicians can order medications at computer terminals not yet "logged out" by the previous physician. This can result in either unintended patients receiving medication or patients not receiving the intended medication. When patients undergo surgery, the CPOE system cancels their previous medications. Physicians must reenter CPOE and reactivate each previously ordered medication. Once again, a reminder to do so may serve to reduce the frequency of such mistakes.

The growing body of research on unintended consequences spurred the American Medical Informatics Association to devote a policy meeting to consider ways to understand and diminish their impact (Bloomrosen et al. 2011). The matter is especially pressing given the increased implementation of health information technologies nationwide including ambulatory care practices that have little experience with health information technologies. The authors outline a series of recommendations, including a need for more cognitively-oriented research to guide study of the causes and mitigation of unintended consequences resulting from health information technology implementations. These changes could facilitate improved management of those consequences, resulting in enhanced performance, patient safety as well as greater user acceptance.

### 4.6.3 External Representations and Information Visualization

To reiterate, internal representations reflect mental states that correspond to the external world. The term external representation refers to any object in the external world which has the potential to be internalized. External representations such as images, graphs, icons, audible sounds, texts with symbols (e.g., letters and numbers), shapes and textures are vital sources of knowledge, means of communication and cultural transmission. The classical model of information-processing cognition viewed external representations as mere

inputs to the mind (Zhang 1997). For example, the visual system would process the information in a display that would serve as input to the cognitive system for further processing (e.g., classifying dermatological lesions), leading to knowledge being retrieved from memory and resulting in a decision or action. These external representations served as a stimulus to be internalized (e.g., memorized) by the system. The hard work is then done by the machinery of the mind, which develops an internal copy of a slice of the external world and stores it as a mental representation. The appropriate internal representation is then retrieved when needed.

This view has changed considerably. Norman (1993) argues that external representations play a critical role in enhancing cognition and intelligent behavior. These durable representations (at least those that are visible) persist in the external world and are continuously available to augment memory, reasoning, and computation. Consider a simple illustration involving multi-digit multiplication with a pencil and paper. First, imagine calculating  $37 \times 93$  without any external aids. Unless you are unusually skilled in such computations, they will exert a reasonably heavy load on working memory in relatively short order. One may have to engage in a serial process of calculation and maintain partial products in working memory (e.g.,  $3 \times 37 = 111$ ). Now consider the use of pencil and paper as illustrated below.

		3	7
	x	9	3
	1	1	1
3	3	3	
3	4	4	1

The individual brings to the task knowledge of the meaning of the symbols (i.e., digits and their place value), arithmetic operators, and addition and multiplication tables (that enable a look-up from memory). The external representations include the positions of the symbols, the partial products of interim calculations and their spatial relations (i.e., rows and columns). The visual representation, by holding partial results outside the mind, extends a person's working

memory (Card et al. 1999). Calculations can rapidly become computationally prohibitive without recourse to cognitive aids. The offloading of computations is a central argument in support of distributed cognition, which is the subject of the next section.

It is widely understood that not all representations are equal for a given task and individual. The representational effect is a well-documented phenomenon in which different representations of a common abstract structure can have a significant effect on reasoning and decision making (Zhang and Norman 1994). For example, different forms of graphical displays can be more or less efficient for certain tasks. A simple example is that Arabic numerals are more efficient for arithmetic (e.g.,  $37 \times 93$ ) than Roman numerals ( $XXXVII \times XCIII$ ) even though the representations or symbols are identical in meaning. Similarly, a digital clock provides an easy read-out for precisely determining the time (Norman 1993). On the other hand, an analog clock provides an interface that enables one to more readily determine time intervals (e.g., elapsed or remaining time) without recourse to calculations. Larkin and Simon (1987) argued that effective displays facilitate problem-solving by allowing users to substitute perceptual operations (i.e., recognition processes) for effortful symbolic operations (e.g., memory retrieval and computationally intensive reasoning) and that displays can reduce the amount of time spent searching for critical information. Research has demonstrated that different forms of graphical representations such as graphs, tables and lists can dramatically change decision-making strategies (Kleinmuntz and Schkade 1993; Scaife and Rogers 1996).

Medical prescriptions are an interesting case in point. Chronic illness affects over 100 million individuals in the United States, many of whom suffer from multiple of these individuals suffer from multiple afflictions and must adhere to complex medication regimens. There are various pill organizers and mnemonic devices designed to promote patient compliance. Although these are helpful, prescriptions written by clinicians are inherently hard for patients to follow. The following prescriptions were given to a patient

following a mild stroke (Day, 1988 reported in Norman 1993).

Inderal	-1 tablet 3 times a day
Lanoxin	-1 tablet every AM
Carafate	-1 tablet before meals and at bedtime
Zantac	-1 tablet every 12 h (twice a day)
Quinaglute	-1 tablet 4 times a day
Coumadin	-1 tablet a day

The physician’s list is concise and presented in a format whereby a pharmacist can readily fill the prescription. However, the organization by medication does not facilitate a patient’s decision of what medications to take at a given time of day. Some computation, memory retrieval (e.g., I took my last dose of Lanoxin 6 h ago) and inference (what medications to bring when leaving one’s home for some duration of hours) are necessary to make such a decision. Day proposed an alternative tabular representation (Table 4.2).

In this matrix representation in Table 4.2, the items can be organized by time of day (columns) and by medication (rows). The patient can simply scan the list by either time of day or medication. This simple change in representation can transform a cognitively taxing task


**Table 4.2** Tabular representation of medications

	Breakfast	Lunch	Dinner	Bedtime
Lanoxin	X			
Inderal	X	X	X	
Quinaglute	X	X	X	X
Carafate	X	X	X	X
Zantac		X		X
Coumadin				X

Adapted from Norman (1988)

into a simpler one that facilitates search (e.g., when do I take Zantac) and computation (e.g., how many pills are taken at dinner time). Tables can support quick and easy lookup and embody a compact and efficient representational device. However, a particular external representation is likely to be effective for some populations of users and not others (Ancker and Kaufman 2007). For example, reading a table requires a certain level of numeracy that is beyond the abilities of certain patients with very basic education. Kaufman and colleagues (2003) characterized the difficulties some older adult patients experienced in dealing with numeric data, especially when represented in tabular form. For example, when reviewing their blood glucose or blood pressure values, several patients appeared to lack an abstract understanding of covariation and how it can be expressed as a functional relationship in a tabular format (i.e., cells and rows) as illustrated in Fig. 4.11. Others had difficulty establishing the correspondence between the values expressed on the interface of their blood pressure monitoring device and mathematical representation in tabular format (systolic/diastolic). The familiar monitoring device provided an easy readout and patients could readily make appropriate inferences (e.g., systolic value is higher than usual) and take appropriate measures. However, when interpreting the same values in a table, certain patients had difficulty recognizing anomalous or abnormal results even when these values were rendered as salient by a color-coding scheme.

The results suggest that even the more literate patients were challenged when drawing inferences over bounded periods of time. They tended to focus on discrete values (i.e., a single reading)



**Fig. 4.11** Mapping values between blood pressure monitor and IDEATel Table



in noting whether it was within their normal or expected range. In at least one case, the problems with the representation seemed to be related to the medium of representation rather than the form of representation. One patient experienced considerable difficulty reading the table on the computer display, but maintained a daily diary with very similar representational properties.

Instructions can be embodied in a range of external representations from text to list of procedures to diagrams exemplifying the steps. Everyday nonexperts are called upon to follow instructions in a variety of application domains (e.g., completing income-tax forms, configuring and using a digital video recorder (DVR), cooking something for the first time, or interpreting medication instructions), where correct processing of information is necessary for proper functioning. The comprehension of written information in such cases frequently involves both quantitative and qualitative reasoning, as well as a minimal familiarity with the application domain. This is nowhere more apparent, and critical, than in the case of over-the-counter pharmaceutical labels, the correct understanding of which often demands that the user translate minimal, quantitative formulas into qualitative, and frequently complex, procedures. Medical errors involving the use of therapeutic drugs are amongst the most frequent.

The calculation of dosages for pharmaceutical instructions can be remarkably complex. Consider the following instructions for an over-the-counter cough syrup.

Each teaspoonful (5 mL) contains 15 mg of dextromethorphan hydrobromide U.S.P., in a palatable yellow, lemon flavored syrup. DOSAGE ADULTS: 1 or 2 teaspoonfuls three or four times daily. DOSAGE CHILDREN: 1 mg/kg of body weight daily in 3 or 4 divided doses.

If you wish to administer medication to a 22 lb child three times a day and wish to determine the dosage, the calculations are as follows:

$$\frac{22\text{lbs}}{2.2\text{lbs/kg}} \times \frac{1\text{mg}}{\text{kg}} \div \frac{15\text{mg}}{\text{tsp}} \div \frac{3\text{doses}}{\text{day}} = \frac{2}{9} \text{tsp} \div \text{dose}$$

Patel, Branch, and Arocha (2002a) studied 48 lay subjects' responses to this problem. The

majority of participants (66.5 %) were unable to correctly calculate the appropriate dosage of cough syrup. Even when calculations were correct, they were unable to estimate the actual amount to administer. There were no significant differences based on cultural or educational background. One of the central problems is that there is a significant mismatch between the designer's conceptual model of the pharmaceutical text and procedures to be followed and the user's mental model of the situation.

*Diagrams* are tools that we use daily in communication, information storage, planning, and problem-solving. Diagrammatic representations are not new devices for communicating ideas. They of course have a long history as tools of science and cultural inventions that augment thinking. For example, the earliest maps, a graphical representation of regions of geographical space, date back thousands of years. The phrase "a picture is worth 10,000 words" is believed to be an ancient Chinese proverb (Larkin and Simon 1987). External representations have always been a vital means for storing, aggregating and communicating patient data. The psychological study of information displays similarly has a long history dating back to Gestalt psychologists beginning around the turn of the twentieth century. They produced a set of laws of pattern perception for describing how we see patterns in visual images (Ware 2003). For example, the **law of proximity** states that visual entities that are close together are perceptually grouped. The **law of symmetry** indicates that symmetric objects are more readily perceived.

Advances in graphical user interfaces afford a wide range of novel external representations. Card and colleagues (1999) define **information visualization** as "the use of computer-supported, interactive, visual representations of abstract data to amplify cognition". Information visualization of medical data is a vigorous area of research and application (Kosara and Miksch 2002; Starren and Johnson 2000). Medical data can include single data elements or more complex data structures. Representations can also be characterized as either numeric (e.g., laboratory data) or non-numeric information such as symptoms and diseases. Visual

representations may be either static or dynamic (changing as additional temporal data become available). EHRs need to include a wide range of data representation types, including both numeric and nonnumeric (Tang and McDonald 2001). EHR data representations are employed in a wide range of clinical, research and administrative tasks by different kinds of users. Medical imaging systems are used for a range of purposes including visual diagnosis (e.g., radiology), assessment and planning, communication, and education and training (Greenes and Brinkley 2001). The purposes of these representations are to display and manipulate digital images to reveal different facets of anatomical structures in either two or three dimensions.

Patient monitoring systems employ static and dynamic representations (e.g., continuously updated observations) for the presentation of the presentation of physiological parameters such as heart rate, respiratory rate, and blood pressure (Gardner and Shabot 2001) (see Chap. 19). As discussed previously, Lin et al. found that that the original display of a PCA device introduced substantial cognitive complexity into the task and impacted performance (Lin et al. 1998). They also demonstrated that redesigning interface in a manner consistent with human factors principles could lead to significantly faster, easier, and more reliable performance.

Information visualization is an area of great importance in bioinformatics research, particularly in relation to genetic sequencing and alignment. The tools and applications are being produced at a very fast pace. Although there is tremendous promise in such modeling systems, we know very little about what constitutes a usable interface for particular tasks. What sorts of competencies or prerequisite skills are necessary to use such representations effectively? There is a significant opportunity for cognitive methods and theories to play an instrumental role in this area.

In general, there have been relatively few cognitive studies characterizing how different kinds of medical data displays impact performance. However, there have been several efforts to develop a typology of medical data representations. Starren and Johnson (2000) proposed a tax-

onomy of data representations. They characterized five major classes of representation types including list, table, graph, icon, and generated text. Each of these data types has distinct measurement properties (e.g., ordinal scales are useful for categorical data) and they are variably suited for different kinds of data, tasks and users. The authors propose some criteria for evaluating the efficacy of a representation including: (1) latency (the amount of time it takes a user to answer a question based on information in the representation), (2) accuracy, and (3) compactness (the relative amount of display space required for the representation). Further research is needed to explore the cognitive consequences of different forms of external medical data representations. For example, what inferences can be more readily gleaned from a tabular representation versus a line chart? How does configuration of objects in a representation affect latency?

At present, computational advances in information visualization have outstripped our understanding of how these resources can be most effectively deployed for particular tasks. However, we are gaining a better understanding of the ways in which external representations can amplify cognition. Card et al. (1999) propose six major ways: (1) by increasing the memory and processing resources available to the users (offloading cognitive work to a display), (2) by reducing the search for information (grouping data strategically), (3) by using visual presentations to enhance the detection of patterns, (4) by using perceptual attention mechanisms for monitoring (e.g., drawing attention to events that require immediate attention), and (5) by encoding information in a manipulable medium (e.g., the user can select different possible views to highlight variables of interest).

#### 4.6.4 Distributed Cognition and Electronic Health Records

In this chapter, we have considered a classical model of information-processing cognition in which mental representations mediate all activity and constitute the central units of analysis. The

analysis emphasizes how an individual formulates internal representations of the external world. To illustrate the point, imagine an expert user of a word processor who can effortlessly negotiate tasks through a combination of key commands and menu selections. The traditional cognitive analysis might account for this skill by suggesting that the user has formed an image or schema of the layout structure of each of eight menus, and retrieves this information from memory each time an action is to be performed. For example, if the goal is to “insert a clip art icon”, the user would simply recall that this is subsumed under pictures that are the ninth item on the “Insert” menu and then execute the action, thereby achieving the goal. However, there are some problems with this model. Mayes, Draper, McGregor, and Koatley (1988) demonstrated that even highly skilled users could not recall the names of menu headers, yet they could routinely make fast and accurate menu selections. The results indicate that many or even most users relied on cues in the display to trigger the right menu selections. This suggests that the display can have a central role in controlling interaction in graphical user interfaces.

As discussed, the conventional information-processing approach has come under criticism for its narrow focus on the rational/cognitive processes of the solitary individual. In the previous section, we consider the relevance of external representations to cognitive activity. The emerging perspective of **distributed cognition** offers a more far-reaching alternative. The distributed view of cognition represents a shift in the study of cognition from being the sole property of the individual to being “stretched” across groups, material artifacts and cultures (Hutchins 1995; Suchman 1987). This viewpoint is increasingly gaining acceptance in cognitive science and human-computer interaction research. In the distributed approach to HCI research, cognition is viewed as a process of coordinating distributed internal (i.e., knowledge) and external representations (e.g., visual displays, manuals). Distributed cognition has two central points of inquiry, one that emphasizes the inherently social and collaborative nature of cognition (e.g., doctors, nurses

and technical support staff in neonatal care unit jointly contributing to a decision process), and one that characterizes the mediating effects of technology or other artifacts on cognition.

The distributed cognition perspective reflects a spectrum of viewpoints on what constitutes the appropriate unit of analysis for the study of cognition. Let us first consider a more radical departure from the classical model of information-processing. Cole and Engestrom (1997) suggest that the natural unit of analysis for the study of human behavior is an activity system, comprising relations among individuals and their proximal, “culturally-organized environments”. A system consisting of individuals, groups of individuals, and technologies can be construed as a single indivisible unit of analysis. Berg is a leading proponent of the sociotechnical point view within the world of medical informatics. He argues that “work practices are conceptualized as networks of people, tools, organizational routines, documents and so forth” (Berg 1999). An emergency ward, outpatient clinic or obstetrics and gynecology department is seen as an inter-related assembly of things (including humans) whose functioning is primarily geared to the delivery of patient care. Berg (1999) goes on to emphasize that the “elements that constitute these networks should then not be seen as discrete, well-circumscribed entities with pre-fixed characteristics (p 89).

In Berg’s view, the study of information systems must reject an approach that segregates individual and collective, human and machine, as well as the social and technical dimensions of IT. Although there are compelling reasons for adapting a strong socially-distributed approach, an individual’s mental representations and external representations are both instrumental tools in cognition (Park et al. 2001; Patel et al. 2002b). This is consistent with a distributed cognition framework that embraces the centrality of external representations as mediators of cognition, but also considers the importance of an individual’s internal representations (Perry 2003).

The mediating role of technology can be evaluated at several levels of analysis from the individual to the organization. Technologies, whether


they be computer-based or an artifact in another medium, transform the ways individuals and groups think. They do not merely augment, enhance or expedite performance, although a given technology may do all of these things. The difference is not merely one of quantitative change, but one that is qualitative in nature.

In a distributed world, what becomes of the individual? We believe it is important to understand how technologies promote enduring changes in individuals Salomon et al. (1991) introduce an important 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 physicians to focus on higher-order thinking skills, such as diagnostic 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. This effect is

illustrated subsequently in the context of the enduring effects of an EHR (see Chap. 12).

We employed a pen-based EHR system, DCI (Dossier of Clinical Information), in several of our studies (see Kushniruk et al. 1996). Using the pen or computer keyboard, physicians can directly enter information into the EHR, such as the patient's chief complaint, past history, history of present illness, laboratory tests, and differential diagnoses. Physicians were encouraged to use the system while collecting data from patients (e.g., during the interview). The DCI system incorporates an extended version of the ICD-9 vocabulary standard (see Chap. 7). The system allows the physician to record information about the patient's differential diagnosis, the ordering of tests, and the prescription of medication. The system also provides supporting reference information in the form of an integrated electronic version of the Merck Manual, drug monographs for medications, and information on laboratory tests. The graphical interface provides a highly structured set of resources for representing a clinical problem as illustrated in Fig. 4.12.

We have studied the use of this EHR in both laboratory-based research (Kushniruk et al. 1996) and in actual clinical settings using cognitive methods (Patel et al. 2000). The laboratory research included a simulated doctor-patient



**Fig. 4.12** Display of a structured electronic medical record with graphical capabilities

interview. We have observed two distinct patterns of EHR usage in the interactive condition, one in which the subject pursues information from the patient predicated on a hypothesis; the second strategy involves the use of the EHR display to provide guidance for asking the patient questions. In the screen-driven strategy, the clinician is using the structured list of findings in the order in which they appear on the display to elicit information. All experienced users of this system appear to have both strategies in their repertoire.

In general, a screen-driven strategy can enhance performance by reducing the cognitive load imposed by information-gathering goals and allow the physician to allocate more cognitive resources toward testing hypotheses and rendering decisions. On the other hand, this strategy can encourage a certain sense of complacency. We observed both effective as well as counter-productive uses of this screen-driven strategy. A more experienced user consciously used the strategy to structure the information-gathering process, whereas a novice user used it less discriminately. In employing this screen-driven strategy, the novice elicited almost all of the relevant findings in a simulated patient encounter. However, she also elicited numerous irrelevant findings and pursued incorrect hypotheses. In this particular case, the subject became too reliant on the technology and had difficulty imposing her own set of working hypotheses to guide the information-gathering and diagnostic-reasoning processes.

The use of a screen-driven strategy is evidence of the ways in which technology transforms clinical cognition, as evidenced in clinicians' patterns of reasoning. Patel and colleagues (2000) extended this line of research to study the cognitive consequences of using the same EHR system in a diabetes clinic. The study considered the following questions (1) How do physicians manage information flow when using an EHR system? (2) What are the differences in the way physicians organize and represent this information using paper-based and EHR systems, and (3) Are there long-term, enduring effects of the use of EHR systems on knowledge representations and clinical reasoning? One study focused on an in-depth characterization of changes in knowledge organi-

zation in a single subject as a function of using the system. The study first compared the contents and structure of patient records produced by the physician using the EHR system and paper-based patient records, using ten pairs of records matched for variables such as patient age and problem type. After having used the system for 6 months, the physician was asked to conduct his next five patient interviews using only hand-written paper records.

The results indicated that the EHRs contained more information relevant to the diagnostic hypotheses. In addition, the structure and content of information was found to correspond to the structured representation of the particular medium. For example, EHRs were found to contain more information about the patient's past medical history, reflecting the query structure of the interface. The paper-based records appear to better preserve the integrity of the time course of the evolution of the patient problem, whereas, this is notably absent from the EHR. Perhaps, the most striking finding is that, after having used the system for 6 months, the structure and content of the physician's paper-based records bore a closer resemblance to the organization of information in the EHR than the paper-based records produced by the physician prior to exposure to the system. This finding is consistent with the enduring *effects of technology* even in absence of the particular system.

Patel et al. (2000) conducted a series of related studies with physicians in the same diabetes clinic. The results of one study replicated and extended the results of the single subject study (reported above) regarding the differential effects of EHRs on paper-based records on represented (recorded) patient information. For example, physicians entered significantly more information about the patient's chief complaint using the EHR. Similarly, physicians represented significantly more information about the history of present illness and review of systems using paper-based records. It's reasonable to assert that such differences are likely to have an impact on clinical decision making. The authors also video-recorded and analyzed 20 doctor-patient computer interactions by 2 physicians with varying levels of expertise.

One of the physicians was an intermediate-level user of the EHR and the other was an expert user. The analysis of the physician-patient interactions revealed that the less expert subject was more strongly influenced by the structure and content of the interface. In particular, he was guided by the order of information on the screen when asking the patient questions and recording the responses. This screen-driven strategy is similar to what we documented in a previous study (Kushniruk et al. 1996). Although the expert user similarly used the EHR system to structure his questions, he was much less bound to the order and sequence of presented information on the EHR screen. This body of research documented both *effects with* and *effects of* technology in the context of EHR use (Salomon et al. 1991). These include effects on knowledge-organization and information-gathering strategies. The authors conclude that given these potentially enduring effects, the use of a particular EHR will almost certainly have a direct effect on medical decision making.

The previously discussed research demonstrates the ways in which information technologies can mediate cognition and even produce enduring changes in how one performs a task. What dimensions of an interface contribute to such changes? What aspects of a display are more likely to facilitate efficient task performance and what aspects are more likely to impede it? Norman (1986) argued that well-designed artifacts could reduce the need for users to remember large amounts of information, whereas poorly designed artifacts increased the knowledge demands on the user and the burden of working memory. In the distributed approach to HCI research, cognition is viewed as a process of coordinating distributed internal and external representations and this in effect constitutes an indivisible information-processing system.

How do artifacts in the external world “participate” in the process of cognition? The ecological approach of perceptual psychologist Gibson was based on the analysis of invariant structures in the environment in relation to perception and action. The concept of *affordance* has gained substantial currency in human computer interaction. It has been used to refer to attributes of objects that

enable individuals to know how to use them (Rogers 2004). When the affordances of an object are perceptually obvious, they render human interactions with objects as effortless. For example, one can often perceive the affordances of a door handle (e.g., afford turning or pushing downwards to open the door) or a water faucet. On the other hand, there are numerous artifacts in which the affordances are less transparent (e.g., door handles that appear to suggest a pulling motion but actually need to be pushed to open the door). External representations constitute affordances in that they can be picked up, analyzed, and processed by perceptual systems alone. According to theories of distributed cognition, most cognitive tasks have an internal and external component (Hutchins 1995), and as a consequence, the problem-solving process involves coordinating information from these representations to produce new information.

One of the appealing features of the distributed cognition paradigm is that it can be used to understand how properties of objects on the screen (e.g., links, buttons) can serve as external representations and reduce cognitive load. The distributed resource model proposed by Wright, Fields, & Harrison (2000) addresses the question of “what information is required to carry out some task and where should it be located: as an interface object or as something that is mentally represented to the user.” The relative difference in the distribution of representations (internal and external) is central to determining the efficacy of a system designed to support a complex task. Wright, Fields, and Harrison (2000) were among the first to develop an explicit model for coding the kinds of resources available in the environment and the ways in which they are embodied on an interface.

Horsky, Kaufman and Patel (2003a) applied the distributed resource model and analysis to a provider order entry system. The goal was to analyze specific order entry tasks such as those involved in admitting a patient to a hospital and then to identify areas of complexity that may impede optimal recorded entries. The research consisted of two component analyses: a cognitive walkthrough evaluation that was modified based

on the distributed resource model and a simulated clinical ordering task performed by seven physicians. The CW analysis revealed that the configuration of resources (e.g., very long menus, complexly configured displays) placed unnecessarily heavy cognitive demands on users, especially those who were new to the system. The resources model was also used to account for patterns of errors produced by clinicians. The authors concluded that the redistribution and reconfiguration of resources may yield guiding principles and design solutions in the development of complex interactive systems.

The distributed cognition framework has proved to be particularly useful in understanding the performance of teams or groups of individuals in a particular work setting (Hutchins 1995). Hazlehurst and colleagues (Hazlehurst et al. 2003, 2007) have drawn on this framework to illuminate the ways in which work in healthcare settings is constituted using shared resources and representations. The *activity system* is the primary explanatory construct. It comprises actors and tools, together with shared understandings among actors that structure interactions in a work setting. The “propagation of representational states through activity systems” is used to explain cognitive behavior and investigate the organization of system and human performance. Following Hazlehurst et al. (2007, p. 540), “a **representational state** is a particular configuration of an information-bearing structure, such as a monitor display, a verbal utterance, or a printed label, that plays some functional role in a process within the system.” The author has used the concept to explain the process of medication ordering in an intensive care unit and the coordinated communications of a surgical team in a heart room.

Kaufman and colleagues (2009) employed the concept of representational states to understand nursing workflow in a complex technology-mediated telehealth setting. They extended the construct by introducing the concept of the “state of the patient” as a kind of representational state that reflects the knowledge about the patient embodied in different individuals and inscribed in different media (e.g., EHRs, displays, paper documents and blood pressure monitors) at a

given point in time. The authors conducted a qualitative study of the ways in which different media and communication practices shaped nursing workflow and patient-centered decision making. The study revealed barriers to the productive use of system technology as well as adaptations that circumvented such limitations. Technologies can be deployed more effectively to establish common ground in clinical communication and can serve to update the state of the patient in a more timely and accurate way.

The framework for distributed cognition is still an emerging one in human-computer interaction. It offers a novel and potentially powerful approach for illuminating the kinds of difficulties users encounter and finding ways to better structure the interaction by redistributing the resources. Distributed cognition analyses may also provide a window into why technologies sometimes fail to reduce errors or even contribute to them.

---

## 4.7 Conclusion

Theories and methods from the cognitive sciences can shed light on a range of issues pertaining to the design and implementation of health information technologies. They can also serve an instrumental role in understanding and enhancing the performance of clinicians and patients as they engage in a range of cognitive tasks related to health. The potential scope of applied cognitive research in biomedical informatics is very broad. Significant inroads have been made in areas such as EHRs and patient safety. However, there are promising areas of future cognitive research that remain largely uncharted. These include understanding how to capitalize on health information technology without compromising patient safety (particularly in providing adequate decision support), understanding how various visual representations/graphical forms mediate reasoning in biomedical informatics and how these representations can be used by patients and health consumers with varying degrees of literacy. These are only a few of the cognitive challenges related to harnessing the potential of cutting-edge technologies in order to improve patient safety.

## Suggested Readings

- Carayon, P. (Ed.). (2007). *Handbook of human factors and ergonomics in health care and patient safety*. Mahwah: Lawrence Erlbaum Associates. A multifaceted introduction to many of the issues related to human factors, healthcare and patient safety.
- Carroll, J. M. (2003). *HCI models, theories, and frameworks: toward a multidisciplinary science*. San Francisco: Morgan Kaufmann. An edited volume on cognitive approaches to HCI.
- Evans, D. A., & Patel, V. L. (1989). *Cognitive science in medicine*. Cambridge, MA: MIT Press. The first (and only) book devoted to cognitive issues in medicine. This multidisciplinary volume contains chapters by many of the leading figures in the field.
- Norman, D. A. (1993). *Things that make us smart: defending human attributes in the age of the machine*. Reading: Addison-Wesley Pub. Co.. This book addresses significant issues in human-computer interaction in a very readable and entertaining fashion.
- 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. This relatively recent article summarizes new directions in decision-making research. The authors articulate a need for alternative paradigms for the study of medical decision making.
- Patel, V. L., Yoskowitz, N. A., Arocha, J. F., & Shortliffe, E. H. (2009). Cognitive and learning sciences in biomedical and health instructional design: a review with lessons for biomedical informatics education. *Journal of Biomedical Informatics*, 42(1), 176–97. A review of learning and cognition with a particular focus on biomedical informatics.
- Preece, J., Rogers, Y., & Sharp, H. (2007). *Interaction design: beyond human-computer interaction* (2nd ed.). West Sussex: Wiley. A very readable and relatively comprehensive introduction to human-computer interaction.

### Questions for Discussion

1. What are some of the assumptions of the distribute cognition framework? What implications does this approach have for the evaluation of electronic health records?
2. Explain the difference between the *effects of technology* and the *effects with technology*. How can each of these effects contribute to improving patient safety and reducing medical errors?

3. Although diagrams and graphical materials can be extremely useful to illustrate quantitative information, they also present challenges for low numeracy patients. Discuss the various considerations that need to be taken into account in the development of effective quantitative representations for lower literacy populations.
4. The use of electronic health records (EHR) has been shown to differentially affect clinical reasoning relative to paper charts. Briefly characterize the effects they have on reasoning, including those that persist after the clinician ceases to use the system. Speculate about the potential impact of EMRs on patient care.
5. A large urban hospital is planning to implement a provider order entry system. You have been asked to advise them on system usability and to study the cognitive effects of the system on performance. Discuss the issues involved and suggests some of the steps you would take to study system usability.
6. Discuss some of the ways in which external representations can amplify cognition. How can the study of information visualization impact the development of representations and tools for biomedical informatics.
7. “When human error is viewed as a cause rather than a consequence, it serves as a cloak for our ignorance” (Henriksen 2008). Discuss the meaning of this quote in the context of studies of patient safety.
8. Koppel and colleagues (2005) documented two categories of errors in clinicians’ use of CPOE systems: 1) Information errors generated by fragmentation of data and 2) human-machine interface flaws. What are the implications of these error types for system design?