Chapter 2
Cognition and Human Computer Interaction in Health and Biomedicine

David R. Kaufman, Thomas G. Kannampallil, and Vimla L. Patel

2.1 Introduction

Do we really need a theory of cognition? What advantages are conferred by a cognitive theory or a collection of theories? How can cognitive theory advance our knowledge as it pertains to the design and use of health information technology? The past 30 years have produced a cumulative body of experiential and practical knowledge about user experience, system design and implementation that provide insights to guide further work. 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 (Patel and Kaufman 2014). Human-computer interaction (HCI) in health care and other domains are at least partly an empirical science where the growing knowledge base can be leveraged as needed. However, practical or empirical knowledge, for example, in the form of case studies is inadequate for producing robust generalizations, or sound design and implementation principles.

D.R. Kaufman (✉)
Arizona State University, Scottsdale, AZ, USA
e-mail: David.Kaufman.1@asu.edu

T.G. Kannampallil
University of Illinois at Chicago, Chicago, IL, USA

The New York Academy of Medicine, New York, NY, USA
e-mail: tgk2@uic.edu

V.L. Patel
Center for Cognitive Studies in Medicine and Public Health, The New York Academy of Medicine, 1216 Fifth Avenue, New York, NY 10029, USA

Columbia University, Weill Cornell College of Medicine, New York, NY, USA

Arizona State University, Scottsdale, AZ, USA
e-mail: vpatel@nyam.org

We argue that there is a need for a theoretical foundation. Of course, theory is a core part of any basic or applied science and is necessary to advance knowledge, to test hypotheses and to discern robust generalizations from the increasingly idiosyncratic field of endeavor.

Cognitive theory has been a central part of HCI since its inception. However, HCI has expanded greatly since its beginning as a discipline focused on a small subset of interactive tasks such as text editing, information retrieval and software programming (Grudin 2008). It is currently a flourishing area of inquiry that covers all manners of interactions with technology from smart phones to ticketing kiosks. Similarly, in health care, HCI research has focused on an enormous range of health information technologies from electronic health record (EHR) systems to consumer fitness devices such as the Fitbit™. In addition, technology is no longer the realm of the solo agent; rather, it is increasingly a team game. This has led to the adaptation of cognitive theories to HCI that stress the importance of the social and/or distributed nature of computing (Rogers 2004).

Rogers (2004, 2012) critiques the rapid pace of theory change. She argues "the paint has barely dried for one theory before a new coat is applied. It makes it difficult for anything to become established and widely used." Although we perceive this to be a legitimate criticism, we must acknowledge the extraordinary diversity in HCI subjects of inquiry. In addition, cognitive theories have endured; however, they have also evolved in response to new sets of circumstances such as the emphasis on real-world research in complex messy settings, on the role of artifacts as mediators of performance and on team cognition.

What role can theory play in HCI research and application? Bederson and Shneiderman (2003) categorize five types of theories that can inform HCI practice:

- Descriptive – providing concepts, terminology, methods and focusing further inquiry;
- Explanatory – elucidating relationships and processes (e.g., explaining why user performance on a given system is suboptimal);
- Predictive – enabling predictions to be made about user performance or of a given system (e.g., predicting increased accuracy or efficiency as a result of a new design);
- Prescriptive – providing guidance for design from high level principles to specific design solutions;
- Generative – seeding novel ideas for design including prototype development and new paradigms of interaction.

Cognitive theories have played an instrumental role in all five categories, although predicting performance across a spectrum of users (e.g., from novice to expert) remains a challenge. In addition, generative theories have begun to play a more central role in HCI design. Although theoretical frameworks such as ethnomethodology, activity theory and ecological psychology, to name a few, have made substantive contributions to the field, this chapter is focused primarily on cognitive theories including classical human information processing, external cognition and distributed cognition.
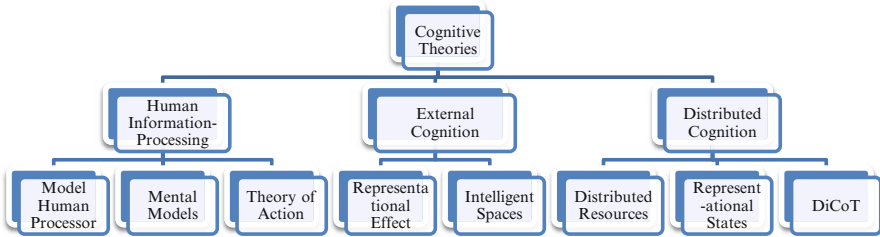
Fig. 2.1 Partial space of frameworks and cognitive theories

In this chapter, we take a historical approach in documenting the evolution of cognitive theories beginning with the early application of information-processing theories and exploring external as well as distributed cognition. Each of these constitutes a family of theories or a framework that embraces core principles, but differs in important respects. A framework is a general pool of constructs for understanding a domain, but it is not sufficiently cohesive or fully realized to constitute a theory (Anderson 1983). The field of HCI as applied to healthcare is remarkably broad in scope and the domain of medicine is characterized by immense complexity and diversity in both tasks and activities (Kannampallil et al. 2011). Specific HCI theories are often limited in scope especially as applied to a rich and complex knowledge domain. Patel and Groen (1992) make an analogous argument for the use of cognitive theories as applied to medical education. Frameworks can provide a theoretical rationale for innovative design concepts and serve to motivate HCI experiments. They can become further differentiated into theories that cover or emphasize a particular facet of interaction (e.g., analyzing teamwork) in the context of a broader framework (e.g., distributed cognition).

We provide a survey of these different theories and illustrate their application with case studies and examples, focusing mostly on issues pertaining to health technology, but also drawing on other domains. This chapter is not intended to be comprehensive or a critical look at the state of the art on HCI in health and biomedicine. Rather, it is written for a diverse audience including those who are new to cognitive science and cognitive psychology. The scope of this chapter is limited with a primary focus on cognitive theories, as they have been applied in healthcare contexts.

A partial space of cognitive theories, as reflected in the chapter, is illustrated in Fig. 2.1. As described, it isn't intended to be exhaustive. It's illustrative of how to conceptualize the theoretical frameworks. It should also be noted that the boundaries between frameworks are somewhat permeable. For example, external and distribute cognition frameworks are co-extensive. However, it serves the purpose of emphasizing the evolution of cognitive theories and highlight specific facets such as the effect of representations on cognition or the social coordination of computer-mediated work. Although the theories within a framework may differ on key issues, the primary difference is in their points of emphasis. In other words, they privilege some aspect as it pertains to cognition and interaction.

2.2   Human Information Processing

A computational theory of mind provides the fundamental underpinning for most contemporary cognitive theories. The basic premise is that much of human cognition can be characterized as a series of operations which reflect computations on mental representations. Early theories and models of human performance were often described in terms of the perceptual and motor activities and assumptions by their structural components (e.g., limits of short-term memory). These were primarily derived from the stimulus-response paradigm, and considered the human as an "information processor." In other words, within this paradigm the human was an information controller, perceiving and responding to activities (Anderson 2005). This approach led to the development of several commonly used models such as Fitts Law (Mackenzie 1992) and the theory of bimanual control (Mackenzie 2003) – that predict performance of human activities in a variety of tasks (e.g., task acquisition, flight controls, and air traffic control). Detailed descriptions of the use of these theories can be found in Chap. 5 of this volume.

With the advent of computers, and more recently significantly interactive environments, there was a need for more integrated information-processing models that accounted for the human-computer interaction (HCI). There were two important requirements: first, the models needed to account for the sequential and integrated actions that evolve during human-computer interactions; second, in addition to the layout and format of the interface, the models also needed to account for the content that was presented on the interfaces (John 2003). In its most general form, the human information processor consists of input, processing and output components (see Fig. 2.2). The input to the processor involves perception of stimuli from the external world; the input/stimuli would be processed by a processor and involves a series of processing stages. Typically, these stages include encoding of the

Fig. 2.2 Input-output model of human information processing. STM refers to short-term memory and LTM is an abbreviation for long-term memory
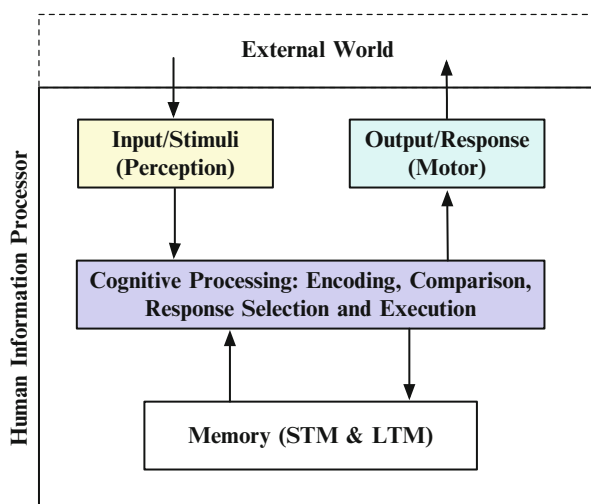
perceived stimuli, comparing and matching it to known mental representations in memory, and selection and execution of an appropriate response. The response is realized through motor actions. For example, consider a clinician's interaction with an EHR interface, where he/she has to select a medication from a dropdown menu. The input component would perceive the dropdown menu from the interface, which would be matched in memory and a click action response would be triggered. This click action would be relayed to the motor components (output), which executes the action by clicking the dropdown menu item. This cycle repeats till the entire task of selecting the medication is completed. In the next sections, we consider core constructs associated with this approach including the model human processor, Norman's theory of action, and mental models.

2.2.1 Model Human Processor

One of the earliest and most commonly described instantiations of a theoretical human information processing system is the Model Human Processor (MHP). MHP can be described as a set of processors, memories and their interactions that operate based on a set of principles (Card et al. 1983). As per MHP, the human mind consists of three interacting processors: perceptual, cognitive and motor. These processors can operate in serial (e.g., pressing a key) or in parallel (e.g., driving a car and listening to radio). Information processing of MHP occurs in cycles. First, the perceptual processor retrieves sensory (visual or audio) information from the external world and is transmitted to the working memory (WM). Once the information is in the WM, information is processed using a *recognize-act* cycle of cognitive processor. During each cycle, contents of WM are connected to actions that are linked to them (from long term memory). These actions, in turn, modify the contents of the WM resulting in a new cycle of actions. MHP can be used to develop an integrated description regarding the psychological effects of human computer interaction performance. While it is considered a significant oversimplification for general users (see applications of the MHP using the GOMS model in Chap. 5), it provided a preliminary mechanism on which much of the human performance modeling research was developed. MHP is useful to predict and compare different interface designs, task performance and learnability of user interfaces. It can be used to develop guidelines for interface design such as spatial layout, response rates and recall. It also provides a significant advantage, as these human performance measures can be determined even without a functional prototype or actual users.

Although the use of MHP approach has not commonly been applied in healthcare contexts, there have been a few noteworthy studies. For example, Saitwal et al. (2010) used the keystroke level model (KLM, an instantiation of the GOMS approach) to compute the time taken, and the number of steps required to complete a set of 14 EHR-based tasks. Using this approach, they characterized the challenges of the user interface and identified opportunities for improvement.

Detailed description of this study and the use of the GOMS approach can be found in Chap. 5.

2.2.2 Norman's Theory of Action

In the mid 1980s, cognitive science was beginning to flourish as a discipline and HCI was viewed as both a test bed for these theories and as a domain of practice. The MHP work was indicative of those efforts. At the same time, microcomputers were becoming increasingly common in homes, work and school. As a result, computers were transitioning from being a tool that was used by experts (i.e., computer scientists and those with high degrees of technical expertise) exclusively to one that was used broadly by individuals in all walks of life. Systems at that point in time were particularly unwieldy and often, extremely difficult to learn. In a seminal paper on cognitive engineering (Norman 1986), Norman sought to craft a theory "to understand the fundamental principles behind human action and performance that are relevant for the development of engineering principles of design" (p 32). A second objective was to devise systems that are "pleasant to use."

A critical insight of the theory is the discrepancy between psychologically expressed goals, and the physical controls and variables of a system. For example, a goal may be to scroll down towards the bottom of a document, and a scroll bar embodies the physical controls to realize such a goal. Shneiderman presented a similar analysis in his theory of direct manipulation (Shneiderman 1982). The key question is how an individual's goals and intentions get expressed as a set of physical actions that transform a virtual system and result in the desired change of state (e.g., reaching the intended section of the document). The Norman model draws on many of the same basic cognitive concepts as the MHP model, but embodies it in a seven stage model of action (Norman 1986), illustrated in Fig. 2.3.

The action cycle begins with a *goal*, for example, retrieving a patient's surgical history. The goal is a generic one independent of any system. In this context, let us presuppose that the clinician has access to paper record as well as those in an EHR. The second stage involves the formation of an *intention*, which in this case might be to retrieve the patient record in an EHR. The intention leads to the *specification of an action* sequence, which may include signing on to the system (which in itself may necessitate several actions), engaging a component system or simply a field that can be used to locate a patient in the database, and entering the patient's identifying information (e.g., last name or medical record number, if it is known). The specification results in *executing an action*, which may necessitate several actions. The system responds in some way or in the case of a failed attempt, may not respond at all. A change in system state may or may not provide a clear indication of the new state or a failure to provide feedback as to why the desired state has not appeared (e.g., system provides no indicators of a wait state or why no response is forthcoming). The perceived system response must then be *interpreted* and *evaluated* to determine whether the goal has been achieved. If the response provided by

Fig. 2.3 Norman's action
cycle
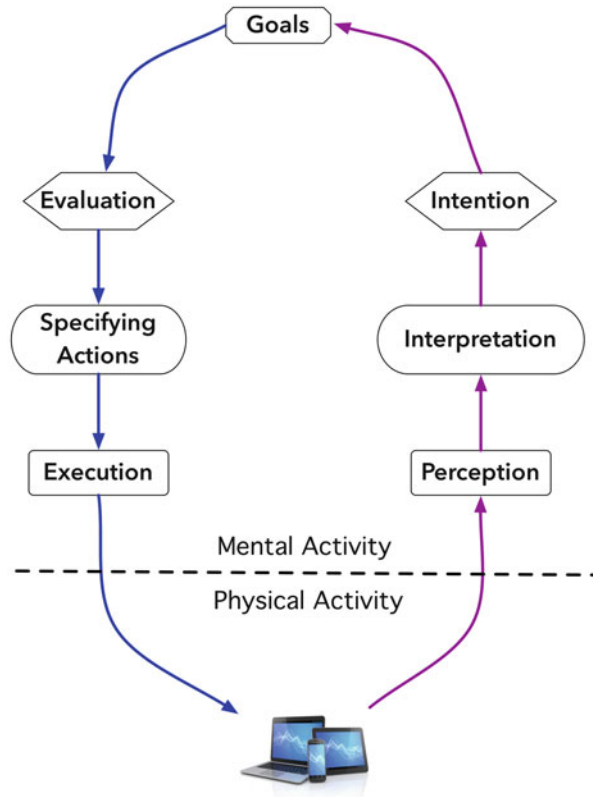
the system is "record not found," that could mean a number of things including that a name was mistyped or the number was incorrectly listed. On the basis of this determination, a next action will be chosen.

Any task of moderate complexity will involve substantial nesting of sub-goals, requiring a series of actions. To an experienced user, the action cycle may appear as a completely transparent and seamless process. However to a less experienced user, the process may breakdown at any of the seven stages. Norman (1986) describes 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. For example, a user may not know the appropriate action sequence or the interface may not provide discernible clues to make such sequences transparent. For instance, a transaction may appear to be complete, but further action is needed to execute the selection process (e.g., pressing enter to accept a transaction).

The *gulf of evaluation* reflects the degree to which the user can make sense of the state of a system and determine how well their expectations have been met. For example, it is sometimes difficult to interpret a state transition and to know whether one has arrived at the correct state or whether the user has chosen an incorrect path.

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 training users to become better attuned to the affordances offered by a system resources. Gulfs can be partially explained by differences in the designer's models and the users' mental models, as discussed in the next section. The designer's model is the conceptual model to be built, based on analysis of the task, requirements, and an understanding of the users' capabilities (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., selecting an item from a menu) will produce predictable and desired outcomes. Graphical user interfaces that involve direct manipulation of screen objects and widgets represent an attempt to reduce the distance between a designer's and user's model (Shneiderman 1982). Obviously, the distance is likely to be more difficult to bridge in a system like an EHR that incorporates a wide range of functions and components that may provide different layouts and forms of interaction.

Norman's theory of action has given rise, or in some cases, reinforced the need for sound design principles. For example, the state of a system should be plainly visible to the user and feedback should be transparent. In illustration, dialog boxes or alert messages can trigger the intention of reminding users to what is possible or needed to complete the task. There is a need to provide good mappings between the actions (e.g., clicking on a tab) and the results of the action as reflected in the state of the system (e.g., providing access to the expected display).

Norman's theory of action informed a great deal of research and design across domains. The seven-stage action theory was used to good effect by Zhang and colleagues in their development of a taxonomy of errors (Zhang et al. 2004). The theory also draws on Reason's categorization of errors as either slips or mistakes (Reason 1992). Slips result from the incorrect execution of a correct action sequence and mistakes are the product of the correct completion of an incorrect action sequence. Slips and mistakes are further categorized into execution errors and evaluation errors. They are further categorized into each of the descriptors that correspond to the Norman's seven stages (e.g., goals, intentions). Zhang et al. (2004) provide the following example of an intention slip: "A nurse intended to enter the rate of infusion using the up–down arrow keys, because this is the technique on the pump she most frequently uses; however, on this pump the arrow keys move the selection region instead of changing the selected number" (p 98). An example of an evaluation/intention slips is that a nurse interprets a yellow flashing light on a device analogically (based on prior knowledge of yellow as a warning) and interprets it as noncritical when it is in fact signaling a critical event. Norman's seven-stage action theory proved to be a useful model for characterizing a wide range of medical error types.

Although theory of action has been very influential in the world of design and research, it also has shortcomings (Sharp et al. 2007). The theory proposes that stages are followed sequentially. However, users do not necessarily proceed in such a sequential manner, especially in a domain such as medicine, which is constituted

by numerous and complex nonlinear tasks. Contemporary GUIs, for example, web-based or app-based systems provide users greater flexibility in achieving the desired state or access the needed information. As discussed in subsequent sections, external representations (e.g., as expressed in text displays or visualizations) offer guidance to the user or even structure their interactions in such a way that a planned action sequence may not be necessary.

2.2.3 Mental Models

Mental models are an important construct in cognitive science and have been widely used in HCI research (Van der Veer and Melguizo 2003). Mental models are an analog-based construct for describing how individuals form internal models of systems. They are employed to answer questions such as "how does it work?" or "what will happen if I make the following move?" "Analog" suggests that the representation explicitly shares some aspect of the structure of the world it represents. For example, one can envision in the mind's eye a set of connected visual images of the succession of ATM screens one has to negotiate to get $200 out of one's checking account or buildings one passes on the way home from a local grocery store. This is in contrast to an abstraction-based form such as propositions or schemas in which the mental structure consists of either the gist, or a summary representation, for example, the procedures needed to complete an ATM transaction. Like other forms of mental representation, mental models are invariably incomplete, imperfect and subject to the processing limitations of the cognitive system (Norman 1983). Mental models can be derived from perception, language or from one's imagination (Payne 2003). Running a model corresponds to a process of mental simulation to generate possible future states of a system from observed or hypothetical state.

The constructs discussed in the prior sections emphasize how the general limits of the human-information processing system (e.g., limits in perception, attention and retrieval from memory) influence performance on a given task in a particular context (Payne 2003). On the other hand, mental models emphasize mental content, namely, knowledge and beliefs. An individual's mental model provides predictive and explanatory capabilities regarding the functions of a particular system. The construct has been used to characterize differences in expertise in a range of knowledge domains such as physics (Payne 2003). Experts have richer and more robust models of a range of phenomena, whereas novices are more prone to imprecision and errors. Mental models 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 downloading a movie from www.amazon.com).

Kaufman et al. (1996) characterized clinician's mental model of the human cardiovascular system (specifically, cardiac output). The study characterized

progressions in understanding of the system as a function of expertise. The research also documented various conceptual flaws in subjects' mental models and how these flaws impacted subjects' predictions and explanations of physiological manifestations (e.g., changes in blood flow in the venous system). In general, mental models are a useful explanatory construct for characterizing errors that are due to problems in understanding and not ones associated with flawed execution of procedures.

Mental models are a particularly useful explanatory device in understanding human-computer interaction (Staggers and Norcio 1993). The premise is that by exploring what users can understand and how they reason about the systems, it is possible to design them in a way that support the acquisition of the appropriate mental model and to reduce errors while performing with them. It is also useful to distinguish between a designer's conceptual model of a given system and a user's mental model (Staggers and Norcio 1993). The wider the gap, the more difficulties individuals will experience in using the system. For example, Kaufman and colleagues (2003) evaluated the usability of a home-based telemedicine system targeting older adults with diabetes. The study documented a substantial gulf between patients' mental models of the system and the designer's intent of how the system should be used. Although most of the participants had a shallow understanding of how such systems worked, there were some who possessed more elaborate mental models, and were better able to negotiate the system to perform a range of tasks including uploading blood glucose values and monitoring one's condition over time.

It is believed that novice users of a system can benefit from instructions that imparts a conceptual model or supports a mental simulation process (i.e., helping the users mentally step through problem states) (Payne 2003). Diagrammatic models of the device or system are often used to support such a learning process. For example, Halasz and Moran (1983) found that such a model was particularly beneficial to students learning to use a programmable calculator. Kieras and Bovair (1984) demonstrated a similar benefit for students learning to master a simple control panel device. They conducted a series of studies contrasting two groups learning to use a device. One group was trained to operate the device through learning the procedures by rote. The second group was trained using a model of how the device works. The model group learned the procedures faster, executed them more rapidly and improvised when necessary, e.g., replacing inefficient procedures with simpler ones. The study provides an illustration of how having a more robust mental model of a system can impact performance. A more advanced model can enable a user to discover alternative ways to achieve the same goal and overcome obstacles.

The construct of mental models fell into disuse in the last couple of decades as theories that emphasized interaction and externalization of representations flourished. However, the construct has resurfaced in recent years as a means to characterize how individuals' conceptualizations differ from representations in systems. For example, Smith and Koppel (2014) take the approach a step further in that they conceptualize three models: the patient's reality, that reality as represented in an EHR and as reflected in a clinician's understanding or mental

model of the problem. Drawing on data from a wide range of sources (e.g., observations and log files) and findings, they constructed "scenarios of misalignment" or misrepresentation including categories such as "IT data too broadly focused" (i.e., lacking precise descriptions). For example, medical problem lists that do not permit sufficient qualification or classification illustrate an example of IT as being too broad or coarse. For instance, clinicians were not able to specify that a stroke resulted from a left-sided cerebrovascular accident. The typology provides a useful basis for IT designers to potentially reduce the gaps, better support users and diminish the potential for unintended consequences.

Shared mental models (SMM) represent an extension of the concept of mental models. The construct is rooted in research on teamwork in areas such as aviation (Orasanu 1990). Clinical care is recognized as a highly collaborative practice and there is a need to develop shared understanding about the processes involved in patient care as well as the evolving conditions of patients that are currently under their care. Breaks in communication among team members are known to be significant contributors to medical errors (Coiera 2000). There are only a few studies that demonstrate a relationship between SMM and clinical performance (Custer et al. 2012). Mamykina and colleagues (2014) investigated the development of SMM in an intensive care unit. The data included observations, audio recorded transcripts of patient handoff (i.e., transfer of patient during shift change) and rounds. In a recent paper, the analysis focused on a single care team including an attending physician, residents, nurses, medical students and physician assistants. The results indicated that the team initially had rather divergent perspectives on how well patients were doing, and the relative success of the treatment. Rounds served as an important coordinating event and the team endeavored to construct shared mental models (i.e., achieving a shared understanding) through an iterative process of resolving discrepancies. There was substantial evidence of change in SMM and in the coordination of patient care over a 3 day period. Whereas conversations on the first day focused on creating basic alignment and making immediate modifications to the care, discussions on the third day focused on understanding of underlying reasons for the situation, and developing a long-term plan more consistent with this collective causal understanding (Mamykina et al. 2014).

As mentioned previously, the concept of mental models has diminished as a construct employed by HCI researchers. One of the reasons is that mental models are not observable and can only be inferred indirectly. However, we believe that it has enduring value as an explanatory device for characterizing how individuals understand a system. The construct is too often used as a synonym for understanding, or for generic mental representation (i.e., with no commitment to the form of the representation). We favor the more specific instantiation of it as a model that can be used to simulate a process and project forward to predict events or outcomes or to explain why a particular outcome occurred. This enables us to develop theories or models for a given domain and then be able to predict and explain variation in performance. This should apply to a wide range of contexts whether the goal is to teach patients with diabetes to understand the basic physiology of their disease or

for clinicians to use a newly implemented EHR. There is also evidence that a model-centric approach to teaching, in which an effort is made to foster an understanding of how a system works, confers some advantages over rote learning approaches to acquire the procedures needed to complete a task (Payne 2003; Gott and Lesgold 2000).

2.3 External Cognition

Internal representations reflect mental states that correspond to the external world. The term external representation refers to any object in the external world that has the potential to be internalized or to be used to augment cognitive processes (without internalizing). External representations such as images, graphs, icons, audible sounds, texts with symbols (e.g., letter 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 that were processed and then internalized (Zhang 1997). The landscape began to change in the early 1990s when new cognitive theories focused on interactivity rather than solely modeling what was assumed to happen inside the head. Rogers (2012) cites Larkin and Simon's (1987) classic paper on "why a diagram may be worth a thousand words" as seminal to researchers in HCI. It offered the first alternative empirical account that focused on how people interact with external representations. The core idea was that cognition can be viewed as the interplay between internal and external representations, rather than only about modeling an individual's mental state and processes. Similar ideas had been put forth by others (Hutchins et al. 1985), but Larkin and Simon provided an explicit computational account that inspired the HCI community (Rogers 2012). Larkin and Simon (1987) made an important distinction between two kinds of external representation: diagrammatic and sentential representations. Although they are informationally equivalent, they are considered to be computationally different. That is, they contain the same information about the problem but the amount of cognitive effort required to come to the solution differs. For example, effective displays facilitate problem solving by allowing users to substitute perceptual operations (i.e., recognition) for effortful cognitive operations (e.g., memory retrieval and computationally-intensive reasoning) and effective displays can reduce the amount of time spent searching for critical information (Patel and Kaufman 2014). On the other hand, cluttered or poorly organized displays may increase the burden.

In the next two sections, we consider two extensions of external cognition, namely, the representational effect and the theory of intelligent spaces.

2.3.1 Representational Effect

The representational effect can be construed as a generalization of Larkin and Simon's (1987) conceptualization of the cognitive impact of external representations (Zhang and Norman 1994). It is well-known that different representations of a common abstract structure can have a significant impact on cognition (Zhang and Norman 1994; Kahneman 2011). For example, different forms of displaying patients' lab values can be more or less efficient for tasks. A display may be oriented to support a quick readout of discrete values or alternatively, one that allows clinicians to discern trends over a period of time. A simple illustration of the effect is that Arabic numerals are more efficient for arithmetic calculations (e.g., 26×92) than Roman numerals (XXVI $\times$ XCII) even though the representations are identical in meaning. Similarly, a digital clock provides a quick readout for precisely determining the time at a glance (Norman 1993). On the other hand, an analog clock enables one to more easily determine time intervals (e.g., elapsed or remaining time) without recourse to mental calculations. Norman (1993) proposed 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. Imagine the cognitive burden of having to do multi-digit multiplication without the use of external aids. Even a pencil and paper will allow you to hold partial results (interim calculations) externally. Calculations can be extremely computationally intensive without recourse to external representations (or memory aids).

Zhang and colleagues (Zhang 1997; Zhang et al.; Zhang and Patel 2006) summarized the following properties of external representations:

- Provide memory aids that can reduce cognitive load
- Provide information that can be directly perceived and used such that minimal processing is needed to explicitly interpret the information
- Support perception so that one can recognize features easily and make inferences directly
- Structure cognitive behavior without cognitive awareness
- Change the nature of a task by generating more efficient action sequences

Several researchers have described the mediating role of information technology on clinical reasoning. For example, Kushniruk et al. (1996) studied how clinicians learned to use an EHR over multiple sessions. They found that as users familiarized themselves with the system, their sequential information-gathering and reasoning strategies were driven by the organization of information on the user interface. In other words, the users followed a "screen-driven" strategy when taking a medical history from a patient. This had both positive consequences in that it promoted a more thorough consideration of the patient history, as well as negative consequences in that the clinician failed to search for findings not available on the display or inconsistent with their operative diagnostic hypothesis. In general, a screen-

driven strategy can enhance performance by reducing the cognitive load imposed by information-gathering goals and enable the physician to allocate more cognitive resources toward patient evaluation (Patel and Kaufman 2014). On the other hand, this strategy can induce a certain sense of complacency or excessive reliance on the display to guide the process.

Similar results were reported by Patel et al. (2000) in a study contrasting the use of EHRs with paper records in a diabetic clinic setting. Physicians entered significantly more information about the patient's chief complaint using the EHR similarly following a screen-driven strategy. Likewise, the structure of the paper records document was such that physicians represented more information about the history of present illness and review of systems using paper-based records. The introduction of an EHR changed information-gathering and documentation strategies, thereby changing the information representation and meaning. The effects of the EHR persisted even after the re-introduction of paper records.

External representations can mediate cognition in a number of ways with both positive and negative impact. The following real-world example was drawn from a study related to a comprehensive causal analysis of a medication dosing error, in which an overdose of Potassium Chloride (KCl) was administered through a commercial computer order entry system (CPOE) in an ICU (Horsky et al. 2005). The authors' detailed analysis included the use of inspection of system logs, interviews with clinicians and a cognitive evaluation of the order-entry system involved. For the purpose of this paper, we highlight one element of the error to illustrate the interplay between technology and user interaction for clinical decision-making. In this case, the system provided screen order-entry forms for medication with intravenous drip and IV bolus orders that were superficially similar, yet required different calculations to estimate the dose. In this case, orders for IV bolus were specified by dose. In contrast, orders for other intravenous drip administration were indicated by duration, rather than by volume of administered fluid as suggested by the order-entry field "Total Volume." The latter referred to the size of the IV bag rather than the total amount of fluid to be delivered, which may exceed the volume indicated. In addition, intravenous fluid orders were not displayed on the medication review screen, further complicating the task of calculating an appropriate KCl bolus for a patient receiving intravenous medications. Calculating the correct infusion dosage was a vitally important task. However, not only did the interface not provide tools to facilitate this process, it also proved to be an obstacle.

It is well documented that IV medication errors commonly result in potentially harmful events (Taxis and Barber 2003; Husch et al. 2005). The configuration of external resources or representations, for example on a visual display, can have a significant impact on how the system facilitates (or alternatively, hinders) cognition. Critical care settings are immensely complex environments and medical error can be the product of a host of factors including workflow and communication (Patel et al. 2014). As discussed in subsequent sections, the organization of displays are just one of several facets that mediate interaction.

2.3.2 *Intelligent Use of Space*

Theories of external cognition tend to emphasize the computational offloading that eases the cognitive burden of a user. However, external representations can also be manipulated by individuals in a variety of ways to facilitate creative thinking as well (Rogers 2012; Zhang and Norman 1994; Kirsh 2005). According to Kirsh, "cognitive processes flow to wherever it is cheaper to perform them. The human 'cognitive operating system' extends to states, structures, and processes outside the mind and body" (Kirsh 2010) (p. 172). For example, one may choose to create a diagram to help interpret a complex sentence and that will alleviate some of the cognitive burden of sense-making. Kirsch draws on a range of examples, in illustration, how people follow a cooking recipe by arranging and re-arranging items (e.g., utensils and ingredients) to coordinate their activities. The central premise is that people interact and create external structure (or representations) because through these interactions, it is easier to process more efficiently and more effectively than by working inside the head alone. In essence, individuals are able to improve their thinking and comprehension by creating and using external representations (Kirsh 2010).

Kirsh (1995) studied how individuals restructured their environments when performing a range of tasks. He found that they constantly rearrange items to track the task state, support memory, predict effects of actions, and so forth. Restructuring often can reduce the cost of visual search, make it easier to notice, identify and remember items, and simplify task representation (Senathirajah et al. 2014a). The theory of intelligent spaces is an extension of this idea. Kirsh classified intelligent uses of space into three categories: (1) arrangements that simplify choice, (2) arrangements that simplify perception (e.g., calling attention to a group of items), and (3) spatial dynamics that simplify mental computation. The theory of intelligent spaces suggests that the idiosyncratic arrangements of individuals including clinicians may serve to simplify inferences or computations. The theory is potentially extensible across a range of domains including health information technology (HIT).

Although EHRs are very elaborate complex systems that support a wide range of functions, they often fail to support the varied needs of healthcare practitioners. Systems often fail to take into consideration the significant variability of medical information needs, which differ according to setting, specialty, role, individual patient and institution (Senathirajah et al. 2014b). In addition, they are not responsive to the highly collaborative nature of the work. In response to these challenges, Senathirajah and colleagues (2014a, b) developed a new model for health information systems, embodied in MedWISE, a widget-based highly configurable EHR platform. MedWISE supports drag/drop user configurations and the sharing of user-created elements such as custom laboratory result panels and user-created interface tabs. It was hypothesized that such a system could afford the clinician greater flexibility and better fit to the tasks they were required to perform. The intelligent spaces theoretical framework informed the design of MedWISE.

In an experiment conducted by Senathirajah et al. (2014b), 13 clinicians used the MedWISE system to review four patient cases. The data included video recordings of clinicians' interactions with the system and the screen layouts they created via the drag/drop capabilities. The focus here was on the creation of spatial layouts. The study documented three strategies which were labeled "opportunistic selection" (rapidly gathering items on the screen and reviewing), structured (organizing the layout categorically) and "dynamic stage" approach. The latter approach involved the user interacting with small groups of widgets at a time, using the space as a staging area to examine a specific concern and then shift to the next. An example of dynamic stage approach was that the clinician kept the index note (initial note) open at the bottom of column 2 (middle column) and stacked the unexamined labs and reports, closed, in column 1 (leftmost column), opened them in column 2 to compare them with the index note, and closed and moved them to column 3. This interaction pattern could reflect examination of specific diagnostic concerns (e.g., ruling out a diagnostic hypothesis). An example of the structured approach is indicated in Fig. 2.4. The clinician has stated that he is keeping the right side as a free space for thinking space, for studies, and for to-do items. A to-do list is at upper

Fig. 2.4 An illustration of a physician using a structured approach in MedWISE

right (in the yellow sticky note), while orienting items including the primary provider clinic note is at left, with lab data down the middle. This reflects a common pattern found of going from left to right with orienting material, data, and then action items. The clinician has grouped labs according to related diagnostic facets, for example, the HbA1c and micro albumin (diabetes-related) are together, and then thyroid-related results (TSH, T3 and T4) are grouped at the bottom of the center column.

The clinicians employed spatial arrangement in ways consistent with theory and research on workplace spatial arrangement (Senathirajah et al. 2014b). This includes assignment of screen regions for particular purposes, juxtaposition of elements to facilitate calculation (e.g., ratios), and grouping elements with common meanings or relevance to the diagnostic facets of the case (e.g., thyroid findings). Clinicians also made deliberate use of the space following a common pattern of left-to-right progression of orienting materials, data, and action items or reflection space. Widget selection was based on an assessment of what information was useful or relevant immediately or likely to be in the near future (as more information is gathered). The study demonstrated how a user-composable EHR in which users have substantial control over how a display is populated and arranged can embody the advantages predicted by the intelligent use of space theory.

The external cognition framework has introduced a set of concepts that has enabled researchers and designers to characterize designs in ways not previously accessible to them (Rogers 2012). As evidenced in the work on MedWISE, it provided a language that framed how people manipulate representations, interact with objects, and organize their space. This provides a basis for designing tools that facilitate different kinds of interaction. It also suggests that there are more and less optimal ways to configure a display for particular tasks and that the impact of such configurations are measurable.

2.4   Distributed Cognition

The external cognition framework seeded important design concepts. It also provides a means to engage in a more rigorous approach to evaluation. The distributed cognition (DCog) approach takes the argument further beyond the internal-external representation divide (Rogers 2012). DCog re-conceptualizes cognitive phenomena in terms of individuals, artifacts, and internal and external representations and their interactions (Rogers 2012). It provides a more extensive account than external cognition. The core approach entails describing a "cognitive system," which involves interactions among people, artifacts they employ, and the environment they are situated in. Hutchins and colleagues proposed a new paradigm for fundamentally rethinking our assumptions about cognition (Hutchins 1995).

DCog represents a shift in the study of cognition from an exclusive focus on the mind of the individual to being "stretched" across groups, material artifacts and cultures (Hutchins 1995; Suchman 1986). This paradigm has gained substantial

currency in HCI research. In the distributed approach, cognition is viewed as a process of coordinating distributed internal (i.e., what's in the mind) and external representations (e.g., visual displays, post-it notes, paper records). Distributed cognition has two focal points of inquiry, one that emphasizes the inherently social and collaborative nature of cognition (e.g., attending physicians, residents, nurses and respiratory therapists in cardiothoracic intensive care unit jointly contributing to a decision process), and one that characterizes the mediating effects of technology (e.g., EHRs, mobile devices apps) or other artifacts on cognition.

Hollan et al. (2000) emphasize that distributed cognition is more than the social distribution of cognitive processes; rather it is a broader conceptualization that includes emergent phenomena in social interactions as well as interactions between people and the structure of their environment. According to Hollan et al., the perspective "highlights three fundamental questions about social interactions: (1) how are the cognitive processes we normally associate with an individual mind implemented in a group of individuals, (2) how do the cognitive properties of groups differ from the cognitive properties of the people who act in those groups, and (3) how are the cognitive properties of individual minds affected by participation in group activities?" (Hollan et al. 2000) (p 177).

DCog is concerned with representational states and the informational flows around the media carrying these representations (Perry 2003). The framework enables researchers to consider all factors relevant to a task, coalescing individuals, the problem and the tools into a single unit of analysis. This makes it a productive means to develop an understanding of how representations act as intermediaries in the dynamically changing and coordinated processes of work activities (Perry 2003).

Hutchins' (1995) seminal analysis of ship navigation of a U.S. navy vessel provided a compelling account of how crews took the ships bearing and how this information was interpreted processed, and transformed across representational states (embodied in media and technology such as ship navigation instruments like the ship's compass and communication among interdependent actors that constitute the ship's crew). The succession of states resulted in the determination of a ships location, progress and how they could be aligned with intended trajectories. The entities operating within the functional system are not viewed from the perspective of the individual, but as a collective (Perry 2003). Both people and artifacts are considered as representational components of the system. As should be clear at this point, external representations are not mere inputs or stimuli to the mind, but play a more instrumental role in cognition.

In the next sections, we review two extensions of DCog including the distributed resource model and the propagation of representational states.

2.4.1 Distributed Resources Model

One of the strengths of the DCog, as applied to HCI, is that it can be used to understand how properties of objects on the screen (e.g., links, menus) can serve as external representations and reduce cognitive load. Wright et al. (2000) proposed a distributed resources model to address the question of the information needed to carry out a task and where it should be located: as an interface object or as knowledge that a user brings to the task. The relative difference in the distribution of representations is pivotal in determining the efficacy of a system designed to support a complex task such as computer provider entry (Horsky et al. 2003). The distributed resources model includes two primary components. The first is a characterization of information structures (i.e., resource types), pertaining to the control of action and the second is a process-oriented description of how these information structures can be used for action (interaction strategies) to complete a task. The information structures can be embodied in any artifact (e.g., paper charts or an EHR). Wright et al. enumerated several of these information structures including plans, goals, history and state. Plans include possible sequence of actions, events, and anticipated states. Goals refer to the desired states the user wants to accomplish. They may be generated internally or emerge from the interaction with the system. History refers to the part of a plan that has already been accomplished. The history of past actions may be maintained in a web browser, for example, as a list of previously visited pages that can be accessed via a drop-down list. State is the current configuration of resources, for example, as represented in the display screen at a given point in time. These are all considered to be resources for action rather than static structures. They can be externalized, manipulated and subjected to evaluation (Wright et al. 2000).

Horsky et al. (2003) employed the distributed resource model to investigate the usability of a CPOE system. The goal was to analyze order-entry tasks and to identify areas of complexity that may impede performance. The research consisted of two component analyses: a cognitive walkthrough evaluation that was modified based on the distributed resource model and an experiment involving a simulated clinical ordering task performed by seven physicians who were experienced users of the CPOE. The walkthrough analysis revealed that the configuration of resources (e.g., very long menus and complexly configured screens) placed an unnecessarily heavy cognitive load on the user. In addition, successful interaction was too often dependent on the recall of system-related knowledge. The resources model was also used to explain patterns of errors produced by clinicians including, selecting an inappropriate order set, omissions and redundant entries. The authors concluded that the reconfiguration of resources may yield guiding principles and design solutions in the development of complex interactive systems (Horsky et al. 2003). In addition, system design that better reflects the constraints of the task (e.g., hospital admission) and domain (e.g., internal medicine) may minimize the need for more robust mental models or extensive system knowledge.

2.4.2 Propagation of Representational States

Horsky et al. conducted a DCog analysis that emphasized the technology-mediating effects of a CPOE interface on clinical performance. Hazlehurst and colleagues (2007) emphasize both the socially-distributed nature and mediating impact of artifacts on communication during cardiac surgery. Towards that end, they employed a cognitive ethnography method to understand how system resources are configured and used for cardiac surgery and to prevent adverse events. DCog focuses on the activity system as the unit of analysis and seeks to understand how properties of this system determine performance (Hutchins 1995; Horsky et al. 2003; Hazlehurst et al. 2007).

Following Hutchins (1995), Hazlehurt views the 'propagation of representational states' through activity systems as explanatory of cognitive behavior and sought to investigate the organizing features of this propagation as an explanation of system and human performance (Hazlehurst et al. 2007). Accordingly, "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 (Hazlehurst et al. 2007) (p 540)". They identified six patterns of communication between surgeon and perfusionist that relate to the functional properties of the activity system. For example, *direction* is a pattern that seeks to transition the activity system to a new state (e.g., administering medications that affect blood coagulation). *Goal sharing* involves creating an expectation of a desired future, but not specifically the action sequence necessary to achieve the target state. These patterns of communication serve to enhance situation awareness, for example, by making the current situation clear and mutually understood.

The distributed cognition approach has been widely used in HCI to examine existing practices and workflow (Rogers 2012). It has also been used to inform the iterative design process by characterizing how the quality and configuration of resources and representations might be transformed and how this change may impact work practices. It is an approach that is inherently well suited to a complex, media-rich and collaborative domain such as medicine. However, a distributed cognitive analysis can be extremely difficult to conduct (requiring substantial specialized knowledge of the analytic approach as well as the knowledge domain), rather complex and very time consuming. In the next section, we describe an approach which endeavors to make the DCog approach more tractable and bring it closer to the design process (Blandford and Furniss 2006).

2.4.3 Distributed Cognition of Teamwork (DiCoT)

DCog's has developed a rather comprehensive and penetrating approach to understanding the different dimensions of human-computer interaction. However, there

is no 'off-the-shelf' methodology for using it in research or as a practitioner (Furniss et al. 2014). According to Rogers, the application of DCog theory and methods are complicated by the fact that there are no set of features to attend to and no checklist or prescribed method to follow (Rogers 2012). In addition, the analysis and abstraction requires a very high level of skill. However, there have been various structured approaches to gathering and analyzing data including the Distributed Resources (DR) Model (Wright et al. 2000) described in a previous section. DiCoT (Distributed Cognition for Teamwork) was developed to provide a structured approach to analyze work systems and teamwork (Furniss et al. 2014; Furniss and Blandford 2006). The approach is informed by theoretical principles from the DCog literature.

The DiCoT framework focuses on developing five interdependent models with different foci: artifacts, physical, information flow, social and evolutionary (Furniss et al. 2014). Each of the models is informed by a set of principles. For example, the artifacts model includes the premise that mediating artifacts are brought into coordination (e.g., paper and electronic health records) in the completion of a task. A second principle is reflected in the fact that we use our environment continuously by "creating scaffolding" to simplify cognitive tasks (Hollan et al. 2000). The physical model refers to the physical organization of work. It is guided by principles such as space and cognition, which states how humans manipulate space towards the facilitation of decision making or problem solving (e.g., grouping objects into categories). This is similar to the intelligent uses of space (Kirsh 1995). Information transformation is one of the principles of information flow. It suggests that transformation occurs when the representation of information changes. As described previously, more effective representations provide better support for reasoning.

DiCoT has been used to analyze complex systems in a range of healthcare contexts including ambulance control room dispatch (Furniss and Blandford 2006) and infusion pump use in intensive care (Rajkomar and Blandford 2012). Emergency medical dispatch is constituted by a team that coordinates the delivery of services (e.g., dispatching an ambulance) to respond to a call for medical assistance. Furniss and Blandford (2006) conducted a study of an EMD team using the DiCoT approach. The focus was on describing the work system, identifying sources of weakness and projecting the likely consequences of a redesign (e.g., what is likely to happen when a centrally available shared display is visible or accessible to each member of the team). On the basis of characterizing systemic weaknesses, they suggested changes to the physical layout that could enhance "cross-boundary working". Their observations revealed a discontinuity between the central ambulance control and the crews in the field. In response, Furniss and Blandford (2006) proposed the use of more flexible communication channels so the crew could be contacted whether they are at a station or are mobile. The multifaceted model enables the researchers to envision a set of consequences to the redesign scheme along a range of dimensions (e.g., information flow). Clinical practitioners and other stakeholders review and comment on the concrete redesign solutions.

The DCog framework, which incorporates a number of interrelated theories, offers the most comprehensive and in our view, the most compelling theoretical approach to explain the technology-mediated and social/collaborative nature of clinical work. Each theory within this framework privileges different aspects of interactions.

2.5 Conclusions

It is reasonable to conclude that we need a theory (or theories) of cognition in the context of HCI and health care. Although we have learned much from empirical studies and applied work, a theoretical framework is needed to account for the broad scope of the field and the complexity that is inherent in the domain of medicine. Without a sound theoretical framework, generalizations would be limited, and principled approaches to design would be largely illusory. In this chapter, we traced the evolution of cognitive theory from the classical information-processing approach to external cognition through distributed cognition. The information-processing approach drew extensively on concepts from cognitive psychology and embraced a computational approach to the study of interaction. The MHP theory (Card et al. 1983) provides insight into cognitive processes and provides a predictive model of behavior, albeit one that is limited in scope. Norman's theory of action (Norman 1986) offers an explanatory account of the challenges involved in using systems. It also offers general prescriptions, for example, emphasizing the importance of quality feedback to the user. The theory of mental models as applied to HCI builds on the idea of gulfs to further explicate the kinds of knowledge needed to productively use a system. It also broadly prescribes how to narrow the divide between designer models and users' mental models. Although these theories are inherently incomplete in their focus on the solitary individual, they continue to be productive as explanatory theories of HCI.

Theories of external cognition expanded the scope of analysis to include a focus on external representations. Several studies have demonstrated how representations mediate cognition and how differential mediation (as reflected in display configurations) can contribute to medical errors. The theory of intelligent spaces (Kirsh 1995) is a generative theory, which seeded concepts that were realized in the design of the MedWISE system. DCog theories are the most encompassing in their focus on both technology-mediated and socially distributed cognition. The theories offer rich descriptive and explanatory accounts of technology use in the medical workplace. Distributed resource theory (Wright et al. 2000) works both as a descriptive theory characterizing the state of affairs and a prescriptive theory that can be used to reconfigure interfaces to alleviate some of the cognitive burden on users. Significant challenges remain in the domain of health information technology. Although cognitive theory cannot provide all of the answers, it remains a powerful tool for advancing knowledge and furthering the scientific enterprise.

Discussion Questions

1. What role can cognitive theory play in HCI research and application? Describe the different kinds of theories that can inform HCI in practice situations.
2. Explain the gulfs of execution and evaluation and how they can be used to inform HCI design.
3. Mental models are an analog-based construct for describing how individuals form internal models of systems. Explain what is meant by analog. How can mental models inform our understanding of the user experience?
4. Describe the meaning and significance of the representational effect. How can it influence the design of visual displays to represent lab results?
5. What implications can one draw from the theory of intelligent spaces? How can it be used to seed design concepts in health care?
6. What are the essential differences between theories of external representation and theories of distributed cognition?

Additional Readings

Carroll, J. M. (2003). *HCI models, theories, and frameworks: Toward a multidisciplinary science*. San Francisco: Morgan Kaufmann.

Patel, V. L., & Kaufman, D. R. (2014). Cognitive science and biomedical informatics. In E. H. Shortliffe & J. J. Cimino (Eds.), *Biomedical informatics: Computer applications in health care and biomedicine* (pp. 133–185). New York: Springer.

Rogers, Y. (2012). HCI theory: Classical, modern, and contemporary. *Synthesis Lectures on Human-Centered Informatics, 5*(2), 1–129.

References

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Cambridge University Press.

Anderson, J. R. (2005). *Cognitive psychology and its implications*. New York: Macmillan.

Bederson, B. B., & Shneiderman, B. (2003). *The craft of information visualization: Readings and reflections*. San Francisco, CA: Morgan Kaufmann.

Blandford, D., & Furniss, D. (2006). DiCoT: A methodology for applying distributed cognition to the design of teamworking systems. In S. W. Gilroy & M. D. Harrison (Eds.), *Interactive systems: Design, specification, and verification* (pp. 26–38). Berlin: Springer.

Card, S. K., Newell, A., & Moran, T. P. (1983). *The psychology of human-computer interaction*. New York: Erlbaum Associates.

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

Custer, J. W., et al. (2012). A qualitative study of expert and team cognition on complex patients in the pediatric intensive care unit. *Pediatric Critical Care Medicine, 13*(3), 278–284.

Furniss, D., & Blandford, A. (2006). Understanding emergency medical dispatch in terms of distributed cognition: A case study. *Ergonomics, 49*(12–13), 1174–1203.

Furniss, D., et al. (2014). Exploring medical device design and use through layers of distributed cognition: How a glucometer is coupled with its context. *Journal of Biomedical Informatics, 53*, 330–341.

Gott, S. P., & Lesgold, A. M. (2000). Competence in the workplace: How cognitive performance models and situated instruction can accelerate skill acquisition. In R. Glaser (Ed.), *Advances in instructional psychology: Educational design and cognitive science* (pp. 239–327). Mahwah: Erlbaum Associates.

Grudin, J. (2008). A moving target: The evolution of human-computer interaction. In *The human-computer interaction handbook–fundamentals, evolving technologies, and emerging applications* (pp. 1–24). Mahwah: Erlbaum Associates.

Halasz, F. G., & Moran, T. P. (1983). Mental models and problem solving in using a calculator. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 212–216). New York: ACM.

Hazlehurst, B., McMullen, C. K., & Gorman, P. N. (2007). Distributed cognition in the heart room: How situation awareness arises from coordinated communications during cardiac surgery. *Journal of Biomedical Informatics, 40*(5), 539–551.

Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI), 7*(2), 174–196.

Horsky, J., et al. (2003). A framework for analyzing the cognitive complexity of computer-assisted clinical ordering. *Journal of Biomedical Informatics, 36*(1), 4–22.

Horsky, J., Kuperman, G. J., & Patel, V. L. (2005). Comprehensive analysis of a medication dosing error related to CPOE. *Journal of the American Medical Informatics Association: JAMIA, 12*(4), 377–382.

Husch, M., et al. (2005). Insights from the sharp end of intravenous medication errors: Implications for infusion pump technology. *Quality & Safety in Health Care, 14*(2), 80–86.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Hutchins, E. L., Hollan, J. D., & Norman, D. A. (1985). Direct manipulation interfaces. *Human-Computer Interaction, 1*(4), 311–338.

John, B. E. (2003). Information processing and skilled behavior. In J. M. Carroll (Ed.), *HCI models, theories and frameworks: Towards a multidisciplinary science* (pp. 55–102). San Francisco: Morgan Kaufmann.

Kahneman, D. (2011). *Thinking, fast and slow*. New York: McMillan Farrar, Straus and Giroux.

Kannampallil, T. G., et al. (2011). Considering complexity in healthcare systems. *Journal of Biomedical Informatics, 44*(6), 943–947.

Kaufman, D. R., Patel, V. L., & Magder, S. A. (1996). The explanatory role of spontaneously generated analogies in reasoning about physiological concepts. *International Journal of Science Education, 18*(3), 369–386.

Kaufman, D. R., et al. (2003). Usability in the real world: Assessing medical information technologies in patients' homes. *Journal of Biomedical Informatics, 36*(1), 45–60.

Kieras, D. E., & Bovair, S. (1984). The role of a mental model in learning to operate a device. *Cognitive Science, 8*(3), 255–273.

Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence, 73*(1), 31–68.

Kirsh, D. (2005). Metacognition, distributed cognition and visual design. In P. Gardenfors & P. Johansson (Eds.), *Cognition, education, and communication technology* (pp. 147–180). London: Routledge.

Kirsh, D. (2010). Thinking with external representations. *AI and Society, 25*, 441–454.

Kushniruk, A. W., et al. (1996). Assessment of a computerized patient record system: A cognitive approach to evaluating medical technology. *MD Computing, 13*(5), 406–415.

Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science, 11*(1), 65–100.

Mackenzie, I. S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction, 7*(1), 91–139.

Mackenzie, I. S. (2003). Motor behavior models for human-computer interaction. In J. M. Carroll (Ed.), *HCI models, theories, and frameworks: Toward a multidisciplinary science*. San Francisco: Morgan Kaufmann.

Mamykina, L., Hum, R. S., & Kaufman, D. R. (2014). Investigating shared mental models in critical care. In V. L. Patel, D. R. Kaufman, & T. Cohen (Eds.), *Cognitive informatics in health and biomedicine* (pp. 291–315). London: Springer.

Norman, D. A. (1983). Some observations on mental models. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 7–14). Hillsdale: Erlbaum Associates.

Norman, D. A. (1986). Cognitive engineering. In D. A. Norman & S. W. Draper (Eds.), *User centered system design* (pp. 31–61). Hillsdale: Erlbaum Associates.

Norman, D. A. (1993). Cognition in the head and in the world: An introduction to the special issue on situated action. *Cognitive Science, 17*(1), 1–6.

Orasanu, J. M. (1990). *Shared mental models and crew decision making* (CSL technical report no. 46). Princeton: Princeton University, Cognitive Sciences Laboratory.

Patel, V. L., & Groen, G. J. (1992). Cognitive framework for clinical reasoning: Application for training and practice. In D. A. Evans & V. L. Patel (Eds.), *Advanced models of cognition for medical training and practice* (pp. 193–212). Heidelberg: Springer-Verlag GmbH & Co. Kg.

Patel, V. L., & Kaufman, D. R. (2014). Cognitive science and biomedical informatics. In E. H. Shortliffe & J. J. Cimino (Eds.), *Biomedical informatics: Computer applications in health care and biomedicine* (pp. 133–185). New York: Springer.

Patel, V. L., et al. (2000). Impact of a computer-based patient record system on data collection, knowledge organization, and reasoning. *Journal of the American Medical Informatics Association: JAMIA, 7*(6), 569–585.

Patel, V. L., Kaufman, D. R., & Cohen, T. (Eds.). (2014). *Cognitive informatics in health and biomedicine: Case studies on critical care, complexity and errors*. London: Springer.

Payne, S. J. (2003). Users' mental models: The very ideas. In J. M. Carroll (Ed.), *HCI models, theories, and frameworks: Toward a multidisciplinary science* (pp. 135–156). San Francisco: Morgan Kaufmann.

Perry, M. (2003). Distributed cognition. In J. M. Carroll (Ed.), *HCI, models, theories, and frameworks: Toward a multidisciplinary science* (pp. 193–223). San Francisco: Morgan Kaufmann.

Rajkomar, A., & Blandford, A. (2012). Understanding infusion administration in the ICU through distributed cognition. *Journal of Biomedical Informatics, 45*(3), 580–590.

Reason, J. (1992). *Human error*. Cambridge: Cambridge University Press.

Rogers, Y. (2004). New theoretical approaches for human-computer interaction. *Annual Review of Information Science and Technology, 38*(1), 87–143.

Rogers, Y. (2012). HCI theory: Classical, modern, and contemporary. *Synthesis Lectures on Human-Centered Informatics, 5*(2), 1–129.

Saitwal, H., et al. (2010). Assessing performance of an electronic health record (EHR) using cognitive task analysis. *International Journal of Medical Informatics, 79*(7), 501–506.

Senathirajah, Y., Bakken, S., & Kaufman, D. (2014a). The clinician in the driver's seat: Part 1 – A drag/drop user-composable electronic health record platform. *Journal of Biomedical Informatics, 52*, 165–176.

Senathirajah, Y., Kaufman, D., & Bakken, S. (2014b). The clinician in the driver's seat: Part 2 – Intelligent uses of space in a drag/drop user-composable electronic health record. *Journal of Biomedical Informatics, 52*, 177–188.

Sharp, H., Rogers, Y., & Preece, J. (2007). *Interaction design: Beyond human-computer interaction*. New York: Wiley.

Shneiderman, B. (1982). The future of interactive systems and the emergence of direct manipulation. *Behaviour & Information Technology, 1*(3), 237–256.

Smith, S. W., & Koppel, R. (2014). Healthcare information technology's relativity problems: A typology of how patients' physical reality, clinicians' mental models, and healthcare

information technology differ. *Journal of the American Medical Informatics Association: JAMIA, 21*(1), 117–131.

Staggers, N., & Norcio, A. F. (1993). Mental models: Concepts for human-computer interaction research. *International Journal of Man-Machine Studies, 38*(4), 587–605.

Suchman, L. (1986). *Plans and situated actions: The problem of human machine interaction*. Cambridge, MA: Cambridge University Press.

Taxis, K., & Barber, N. (2003). Causes of intravenous medication errors: An ethnographic study. *Quality & Safety in Health Care, 12*(5), 343–347.

Van der Veer, G. C., & Melguizo, P. (2003). Mental models. In A. Sears & J. A. Jacko (Eds.), *The human-computer interaction handbook: Fundamentals, evolving technologies, and emerging applications* (pp. 58–80). Mahwah: Erlbaum Associates.

White, B. Y., & Frederiksen, J. R. (1990). Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence, 42*(1), 99–157.

Wright, P. C., Fields, R. E., & Harrison, M. D. (2000). Analyzing human-computer interaction as distributed cognition: The resources model. *Human-Computer Interaction, 15*(1), 1–41.

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

Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science, 18*(1), 87–122.

Zhang, J., & Patel, V. L. (2006). Distributed cognition, representation, and affordance. *Pragmatics & Cognition, 14*(2), 333–341.

Zhang, J., et al. (2004). A cognitive taxonomy of medical error. *Journal of Biomedical Informatics, 37*(3), 193–204.

Zhang, J., Johnson, K. A., Malin, J., & Smith, J. W. (2002). Human-centered information visualization. Paper presented at *International workshop on dynamic visualization and learning*, Tubingen, Germany, July 18–19, 2002.