Cognitive Modeling for Cognitive Engineering

Wayne D. Gray

1 Introduction

Cognitive engineering is the application of cognitive science theories to human factors practice. As this description suggests, there are strong symbioses between cognitive engineering and cognitive science, but there are also strong differences.

Symbiosis implies a mutual influence, and the history of cognitive engineering supports this characterization in two key areas: the development of cognitive theory and the development of computational modeling software. For theory development, a stringent test of our understanding of cognitive processes is whether we can apply our knowledge to real-world problems. The degree to which we succeed at this task is the degree to which we have developed robust and powerful theories. The degree to which we fail at this task is the degree to which more research and stronger theories are required (Gray, Schoelles, & Myers, 2004).

The development of the production-system-based architectures most strongly associated with cognitive engineering [ACT-R (Anderson, 1993), EPIC (Kieras & Meyer, 1997), and Soar (Newell, 1990)] was motivated by the desire to explore basic cognitive processes. However, each has been strongly influenced by a formalism for cognitive task analysis that was developed explicitly for the application of cognitive science to human-computer interaction (Card, Moran, & Newell, 1980a, 1980b, 1983). Indeed, it can be argued that the modern form of ACT-R (Anderson et al., 2004) and the development of EPIC (Kieras & Meyer, 1997) with their modules that run in parallel owes a great intellectual debt to the development of CPM-GOMS (Gray & Boehm-Davis, 2000; John, 1988, 1993). It is definitely the case that the potential of these architectures for application has long been recognized (Elkind, Card, Hochberg, & Huey, 1989; Pew, 2007; Pew & Mavor, 1998) and that much recent development of these basic research architectures has been funded at least partly because of their potential in tutoring systems (S. F. Chipman, personal communication, 2007-04-02), human-computer interaction (Chipman & Kieras, 2004; Freed, Matessa, Remington, & Vera, 2003; Williams, 2000), or human-system integration (Gluck & Pew, 2005; Gray & Pew, 2004).

On the other hand, the engineering enterprise of building systems that are in some way directly relevant to real world problems is fundamentally different from the basic research enterprise of developing or elaborating cognitive theory. Cognitive science and cognitive engineering can be viewed as differing along five dimensions. Although these differences do not imply a dichotomy, they can be viewed as capturing some of the characteristic differences of these two endeavors.

First is the nature of the problems picked. As an applied discipline, the problems addressed by cognitive engineering are often not picked by the researcher, but are defined for the researcher in terms of safety, workload, design, operational need, or financial impact.

Second is the amount of prior study of the task and task domain. Many of our best models of cognitive theory rest on years of exploring a small number of experimental paradigms within a well-specific domain. Great examples of this would be models of reasoning (Johnson-Laird, 1993; Rips, 1994), models of category learning (Love, Medin, &

Gureckis, 2004; Nosofsky & Palmeri, 1997; Shepard, Hovland, & Jenkins, 1961), as well as models of memory retrieval (Anderson & Schooler, 1991; Hintzman, 2005). In contrast, many computational models for cognitive engineering tend to be first-generation attempts in that little, if any, prior empirical or modeling work exists. Two examples that are discussed in this chapter are Byrne and Kirlik's (2005) work on modeling the taxiing behavior of commercial airline pilots and Gluck's work on modeling uninhabited air vehicle operators (Gluck, Ball, & Krusmark, 2007).

Third, many but not all computational models for cognitive engineering entail domain-specific expertise. This characterization applies to both the development of tutoring systems for the training of novices as well as to the modeling of expert performance. It is definitely the case that much has been learned about basic cognitive processes by studying the acquisition or execution of expertise (Chi, Feltovich, & Glaser, 1981). It is also the case that there is a vast middle ground of educational research in which the distinction between basic versus domain-specific work is often blurred (Anderson, Conrad, & Corbett, 1989; Corbett & Anderson, 1988; Singley & Anderson, 1989). However, at the further extreme are the attempts to model rare forms of expertise, such as that possessed by Submarine Commanders (Ehret, Gray, & Kirschenbaum, 2000; Gray, Kirschenbaum, & Ehret, 1997; Gray & Kirschenbaum, 2000; Kirschenbaum & Gray, 2000), uninhabited air vehicle (UAV) operators (Gluck et al., 2007), or airline pilots (Byrne & Kirlik, 2005). Although, arguably, insights and progress into basic research issues have emerged from these studies, it is undoubtedly true that the motivation and funding to study and the particular expertise of such small populations[1] stems from the need to solve very important applied problems.

Fourth, computational modeling for cognitive engineering operates in an arena where the demand for answers is more important than the demand for understanding. Newell warned us about such arenas (Newell & Card, 1985); if another discipline can reduce human errors, increase productivity, and in general augment cognition then who cares if those advances rely on an in-depth understanding of the human cognitive architecture? The issue for cognitive science is one of relevance.[2]

Fifth, whereas many of our best cognitive science models focus on the distilled essence of a cognitive functionality such as memory or categorization, cognitive engineering models are called on to predict performance in task environments that entail many cognitive functionalities. Hence, the particular challenge of computational modeling for cognitive engineering is to model not just the pieces but also the control of an integrated cognitive system (Gray, 2007b).

These characteristic differences between basic and applied computational cognitive modeling are not meant as dichotomies, but rather to illustrate the different sets of challenges faced by cognitive engineering. To some degree these challenges can be seen as challenges for the basic science; especially the need for cognitive engineering to

[1] For example, the active duty population of Submarine Commanders is estimated to be less than 100.

[2] A reviewer for this chapter proposed astronomy as an example in which public funding continues to flow in the absence of any immediate relevance to the human condition. It is not clear, however, that this example actually makes that case. Indeed, the evidence suggests just the opposite. Astronomy is the smallest field that NSF tracks in their surveys of doctoral scientists and engineers in the US (Tsapogas, 2003). Staying just within the NSF defined category of Physical Sciences in 2003 there were 4,280 living Astronomers in the US (including retired, unemployed, and employed) as compared to 69,460 Chemists (excluding biochemistry), 20,220 Earth Scientists, and 40,440 Physicists. Astronomers are fond of pointing out that expensive space programs such as the Shuttle is not astronomy and that the bulk of the money in expensive "Big Science" program such as the Hubble Telescope and deep space probes goes to engineering not astronomy.

model the control of integrated cognitive systems (the last item on my list). Unfortunately, neither the list nor the efforts that instantiate it are tidy.

The next section reviews the seminal work of Card, Moran, and Newell (Card et al., 1983) from the modern perspective. We then jump to the 2000s to discuss the issues and applications of cognitive engineering, first for the broad category of complex systems and then for the classic area of human-computer interaction, with a focus on human interaction with quantitative information, that is, visual analytics[3]. The chapter ends with a summary and discussion of cognitive engineering.

2 Initial Approaches to Cognitive Modeling for Cognitive Engineering

Attempts to apply computational and mathematical modeling techniques to human factors issues have a long and detailed history. Unfortunately, we cannot review that history here; however, we can do the next best thing and point the reader to Dick Pew's (Pew, 2007) very personal history of human performance modeling from the 50's on. In this section, we pick up the cognitive science side of the story with Card, Moran, and Newell's seminal GOMS[4] framework for applying the information-processing approach to developing cognitive task analysis.

Before the cognitive revolution and, arguably, continuing today, most researchers studying cognitive human behavior were trained in experimental psychology. This tradition focuses on teasing and torturing secrets from nature by tightly controlled studies in which small manipulations are made, and humans perform many nearly identical trials. People with this background and training often cannot conceive how someone could possibly study, let alone model, something as complex as VCR programming (Gray, 2000), driving (Salvucci, 2006), the influence of the layout of a graph on performance (Peebles & Cheng, 2003), information search on the World Wide Web (Blackmon, Kitajima, & Polson, 2005; Kaur & Hornof, 2005; Pirolli, 2005), or air traffic control (ATC) issues (Byrne & Kirlik, 2005).

Although the study of such issues is complex and demanding, it is made possible by an open secret long exploited by the human factors community (Kirwan & Ainsworth, 1992) and long recognized by cognitive science (see Simon, 1996, ch 8 – The Architecture of Complexity, and his discussion therein of "near decomposability"); namely, that most any human behavior that extends in time longer than a few minutes can be conceived of as a hierarchical series of tasks, subtasks, and subsubtasks. The structure of this hierarchy is, for the most part, determined by the nature of the task and task environment and less so by the human operator. Rather than having to deal with hours of behavior as one unit, the human factors analyst can break the behavior down to the level required by the goals of the analysis.

For the human factors professional, this *task analysis* approach works well for designing complex industrial operations in which the lowest unit of analysis is the human operator as well as designing procedures for individual

---

[3] Where visual analytics is defined as the visual representation of quantitative data (Thomas & Cook, 2005; Wong & Thomas, 2004). The human-computer interaction interest in visual analytics lies in the building of interfaces that support the search and representation of massive quantities of quantitative data.

[4] A GOMS task analysis analyzes human behavior in terms of its goals, the operators needed to accomplish the goals, sequences of operators and subgoals that constitute methods for accomplishing a goal, and selection rules for choosing a method when alternative methods for accomplishing the same goal exist.

humans in which each low-level task requires minutes or hours to perform (Kirwan & Ainsworth, 1992; Shepherd, 1998, 2001). For the cognitive scientist interested in interactive behavior, arguably, the job is even easier. Although behaviors may extend indefinitely in time, most interactive behavior results from and results in changes to the task environment. For the pilot of an F-16, the driver of a car, or the user of a VCR, the paradigm comes down to (a) do something, (b) evaluate change in the world, and (c) return to (a). Although interactive behavior is complex, the complexity lies not in planning and executing a long sequence of behavior, but (a) evaluating the current state of the task environment, (b) deciding what can be done "now" that will advance the user's goals given the current state of the task environment, (c) evaluating the strategies available to the human operator for accomplishing the current, momentary goal, and (d) executing *(c)* to accomplish *(b)*. The key to this interactive cycle is the *unit task*.

2.1 The Unit Task as a Control Construct for Human Interactive Behavior

Card, Moran, and Newell's conceptual breakthrough was that even tasks which lasted only minutes were composed from a series of smaller "*unit tasks* within which behavior is highly integrated and between which dependencies are minimal. This quasi-independence of unit tasks means that their effects are approximately additive" (Card et al., 1983, p. 313). The "unit task is fundamentally a control construct, not a task construct" (Card et al., 1983, p. 386). The unit task is not given by the task environment, but results from the interaction of the task structure with the control problems faced by the user.

The prototypical example of a unit task (from Chapter 11 in Card et al., 1983) is the structure imposed by a typist on transcription typing. The physical task environment for transcription typing consists of the dictated speech, a word processor, and a foot pedal that controls how much of a recording is played back. As speech is typically much faster than skilled typing the basic problem faced by the typist is how much of the recording to listen to before shutting it off. The efficient typist listens while typing, and the longer he or she listens, the greater the lag between what they are hearing and what they are typing. At some point, the typist shuts off the recording and continues to type until s/he can remember no more of the recording with certainty. With some experience with the particular speaker and maybe with the particular topic, a skilled transcription typist will minimize the amount of rewind and replay, and maximize the amount typed per unit task. This chopping up of the physical task environment into unit tasks reflects a control process that adjusts performance to the characteristics of the task (the speed of dictation and clarity of speech), the typist's general typing skill (number of words per minute), and to the typist's cognitive, perceptual, and motor limits.

2.2 The Path from Unit Tasks through Interactive Routines to Embodiment

A typical GOMS unit task is shown in Table 1. This sample unit task is one of approximately 20 needed to model Lovett's (Lovett & Anderson, 1996) building sticks task, a simple game whose objective is to match the length of a target stick by building a new stick from pieces of various sizes (a dry analogue to the better-known water jug problem). This unit task would be invoked to subtract length from the built stick when it is larger than the target stick. This example shows that each line or statement has an execution overhead (Stmt Time) of 0.1 second. There are three types of operators (Ops) used. P is a point operator that is assumed to have a time of 1.1 seconds. BB is a button click (up and down) with the duration of 0.2 sec. M is a mental operator with the duration of 1.2 sec. The entire method for accomplishing this unit task lasts 5.8 sec.

Table 1: Example Unit Task for the “Building Sticks Task” using Natural GOMS Language (NGOMSL) format (Kieras, 1997).

	Step	Description	Stmt Time	Op	# Ops	Op Time	Total Time
Method for goal: Subtract stick<position>			0.1				0.1
	Step 1	Point to stick<position>	0.1	P	1	1.1	1.2
	Step 2	Mouse click stick<position>	0.1	BB	1	0.2	0.3
	Step 3	Confirm: Stick is now black	0.1	M	1	1.2	1.3
	Step 4	Point to inside of "your stick"	0.1	P	1	1.1	1.2
	Step 6	Click mouse	0.1	BB	1	0.2	0.3
	Step 7	Confirm: Change in stick size	0.1	M	1	1.2	1.3
	Step 8	Return with goal accomplished	0.1				0.1
							5.8

Stmt Time = statement time; Op = operator; P = point operator, BB = button click; M = mental operator. All times are in seconds.

As suggested by the table, the NGOMSL format reduces all operators to one of a small set. The duration of each operator is based on empirical data, mathematical descriptions of behavior such as Fitts’ Law or Hicks Law, and so on. Much of what goes into an NGOMSL analysis comes out of Card, et al.’s (1983, Chapter 2) Model Human Processor. That chapter summarizes many important regularities gleaned from experimental psychology but cast into a form that could be used by human factors analysts.

GOMS was intended as a tool for cognitive engineering. Hence, whereas each line of the NGOMSL analysis could be made more precise and more tailored to, say, the exact distance moved, a large motivation for GOMS was to derive engineering-style approximations for predicting human behavior. However, for some applied purposes the grain size of GOMS analyses in Table 1 is too gross. Indeed, to model transcription typing John (1996) had to develop a version of GOMS that went below the grain size of normal GOMS. John (1988) represented the dependencies between cognitive, perceptual, and motor operations during task performance (see Figure 1) in an activity network formalism (Schweickert, Fisher, & Proctor, 2003) that allowed the computation of critical paths. This version of GOMS is called CPM-GOMS, where the CPM has a double meaning as both critical path method and cognitive, perceptual, and motor operations.

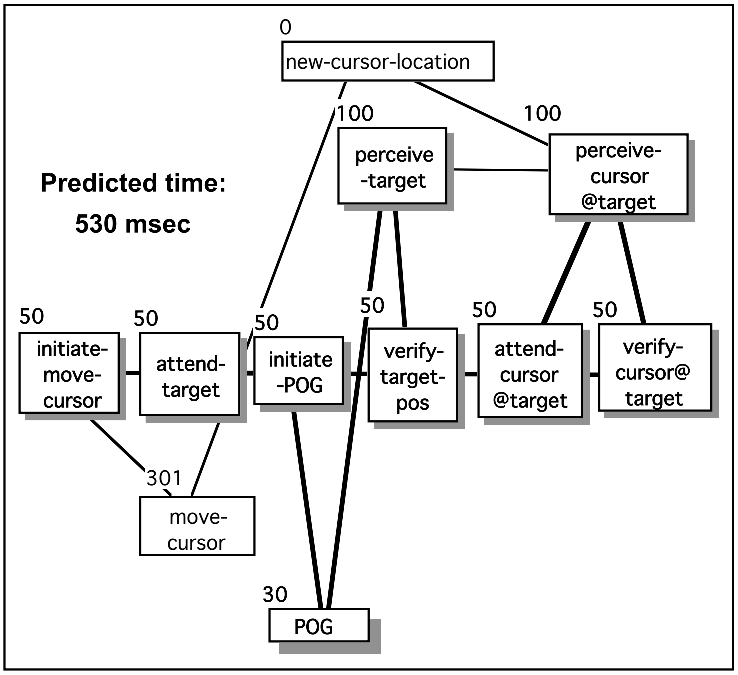

Figure 1: A CPM-GOMS model of an interactive routine. This interactive routine could be instantiated as Step 1 and Step 4 in Table 1. It shows the cognitive, perceptual, and motor operations required to move a mouse to a predetermined location on a computer screen. Total predicted time is 530 ms. The middle row shows cognitive operators with a default execution time of 50 ms each. Above that line are the perceptual operators, and below it are the motor operators. The flow of operators is from left-to-right with connecting lines indicating dependencies. Within an operator type, the dependencies are sequential. However, between operator types the dependencies may be parallel. The numbers above each operator indicate the time, in milliseconds, for that operator to execute. Time is accumulated from left-to-right along the critical path. [The critical path is indicated by bold lines connecting the shadowed boxes.] loc = location; POG = point of gaze. See Gray and Boehm-Davis (2000) for more detailed information.

The power of this representation received a boost from its ability to predict performance times in a very prominent field test of two workstations for telephone Toll and Assistance Operators (TAOs) (Gray, John, & Atwood, 1993). Not only did CPM-GOMS models predict the counterintuitive finding that TAOs using a proposed new workstation would perform more slowly than those who used the older workstations, but after a field trial confirmed this prediction, the models provided a diagnosis in terms of the procedures imposed by workstations on the TAO as to how and why newer, faster technology could perform more slowly than older technology.

2.3 The Legacy of Card, Moran, and Newell

Representations have a power to make certain things obvious, and the GOMS and CPM-GOMS representations did so in several ways. First was the basic insight offered by the unit task; namely, that functional units of behavior resulted from an interaction between the task being performed, detailed elements of the design of the task environment, and limits of human cognitive, perceptual, and motor operations. Second, the notation of CPM-GOMS made it very clear that all human behavior was embodied behavior. Indeed, the mechanistic representations of CPM-GOMS were very compatible with the views of embodiment expressed by modelers such as Ballard (Ballard, Hayhoe, & Pelz, 1995; Ballard, Hayhoe, Pook, & Rao, 1997; Ballard & Sprague, 2007) and Kieras (Kieras & Meyer, 1997) and, at the same time, completely side-stepped the extreme, philosophical claims that are sometimes attached to this concept (Wilson, 2002). Third, whereas standard GOMS and NGOMSL (Kieras, 1997) emphasized control of cognition, CPM-GOMS provided a representation that showed that this control was far from linear, but entailed a complex interleaving of various parallel activities.

Whether as part of the Zeitgeist or as a driving force, the 90's saw many of the insights of CPM-GOMS become standard among modelers. Kieras and Myers built a new cognitive architecture, EPIC (Kieras & Meyer, 1997), by expanding Kieras' parsimonious production system (Bovair, Kieras, & Polson, 1990; Kieras & Bovair, 1986) to include separate modules for motor movement, eye movements, and so on. ACT-R (Anderson, 1993) flirted with the addition of a module for visual attention (Anderson, Matessa, & Lebiere, 1997), experimented with a graft of EPIC's modules (Byrne & Anderson, 1998), and completely restructured itself so that all cognitive activity (not simply that which required interactive behavior) entailed puts and calls to a modular mind (Anderson et al., 2004). During the same period, Ballard's notions of embodiment (Ballard et al., 1997) took literal form in Walter – a virtual human who could follow a path while avoiding obstacles, picking up trash, and stopping to check traffic before he crossed the street (Ballard & Sprague, 2007).

GOMS and the concept of the unit task were conceived as tools to develop "an engineering-style theory of how the user interacts with the computer" (Newell & Card, 1985) in an effort to "harden" the science base for cognitive engineering. During the 90's the attention of the basic research cognitive science community turned to issues of control of cognition, exactly those issues that were highlighted first by the unit task and then by CPM-GOMS. Although by no means complete, by the turn of the 21st century, the theoretical tools and modeling techniques were in place to accelerate progress on the cognitive engineering agenda.

3 Issues and Applications of Computational Modeling for Cognitive Engineering

To illustrate the differences between contemporary cognitive engineering and contemporary cognitive science, a very selective review of two areas of recent research is provided. The first is the broad area of Complex Systems. Work in this area has the typical human factors character of one team of researchers working on one applied problem. The second area is the human-computer interaction topic of visual analytics – how best to design computer interfaces (where "interface" includes "interactive procedures") to facilitate the search and understanding of abstract data presented in visual form. Work in this area resembles the type of distributed research activity familiar to researchers in areas such as visual attention, memory, or categorization.

The position adopted in this section is that the contribution of cognitive engineering is in solving applied problems and in identifying gaps in the underlying cognitive theory. To address this first point, more time than the reader might expect is spent on explaining the domain as well as explaining the nature of the problem that is being solved. To address the second point, details of the model are not discussed. As most cognitive engineering relies on modeling techniques developed elsewhere, those interested in these details may turn to the original sources. Rather, the focus here will be on identifying the special problems that the modelers faced, theoretical mechanisms that contributed to the success of the applied model, and identifying the gaps in cognitive theory that the applied model revealed.

3.1 Complex Systems

A major goal of cognitive engineering is to design high-fidelity models of the demands on human cognitive, perceptual, and action resources during the operation of complex, technological systems. The level of analysis of cognitive engineering is much like that of cognitive science. However, a characteristic difference is that in the typical laboratory study for, say, memory or visual search, the task being studied is the main task being performed. In contrast, cognitive engineering tends not to focus on the main task per se, but on a piece of the main task. So in the arena of commercial aviation, the focus is not on the successful operation of an Air Traffic Control (ATC) system or even on the take-off, flight, and successful landing of individual flights. Much more typically, the focus would be on a small portion of the flight, such as why pilots get lost while taxiing to their gate after landing. Likewise, for the case of driving a car on a crowded highway, cognitive engineering is not concerned with variables such as the number of drivers on the road, miles driven, number of accidents, and so on. Rather, cognitive engineering would focus on basic questions concerning how best to design instrument panels and what sorts of guidelines should be given to manufacturers regarding the design of in-vehicle systems such as car radios, navigation systems, and cell phones (Green, 1999, 2002).

3.1.1 An Operational Challenge: The Predator Uninhabited Air Vehicle

An important challenge for cognitive engineering is the design of new systems, especially those that create new roles for human operators. One such system is UAV. UAVs are increasingly used by the military and intelligence agencies in place of humanly piloted aircraft for a variety of missions. There is also some thought that in the foreseeable future, UAVs may replace some portion of human-piloted commercial aviation (Gluck et al., 2005).

Remotely piloting a slow-moving aircraft while searching for ground targets is a mission that is difficult for even experienced Air Force pilots. A complete model that could take off, perform missions, and return safely would entail the detailed integration of most, if not all, functional subsystems studied by cognitive scientists today as well as raising challenging issues in the control of integrated cognitive systems. Such a complete system is beyond our current state-of-the-art. However, partial systems can be useful in determining limits of human performance and identifying strategies that work. Such partial models have been built by Gluck and colleagues (Gluck et al., 2007) to study the challenges to the human pilot in three routine UAV maneuvers. These researchers modeled two alternative sets of strategies and were able to show that one set would not meet the performance demands of the UAV, whereas the other set would. Close analysis of human performance data suggested that the best human pilots used the strategies incorporated into the best performing model.

Unlike tasks such as simple decision making or categorization, a key challenge to the modelers was obtaining access to an adequate simulation of the pilot's task environment, namely the aerodynamics of a UAV in flight. The flight dynamics of a UAV are very different from those of manned vehicles, and understanding these dynamics presents a challenge for even experienced Air Force pilots. Given that the UAV is traveling at such and such an altitude and speed, what needs to be done to turn it right by 25° while descending by 1,000 feet and slowing by 50 mph within a given period of time without stalling? Computing the effect of such changes on UAV flight in such a dynamic task environment is not trivial and, indeed, is a significant aerodynamic engineering effort. An additional problem not faced by basic theory cognitive science modelers is that UAVs are artifacts that are being constantly upgraded.

Indeed, the most recent UAVs have very different flight dynamics than the UAV used by Gluck and associates. This interesting program of research has slowed to a halt, as the flight characteristics of the new UAVs are sufficiently different from the old UAVs to make cognitive modeling impossible without a new aerodynamic model.

3.1.2 Runway Incursions or Getting Lost While Driving an 870,000 Pound Jetliner to the Gate

Making a wrong turn while driving is frustrating. Making a wrong turn after landing your passenger jet and trying to taxi to your gate is a *runway incursion* which is considered by Air Traffic Control as "creating a collision hazard with another aircraft taking off or landing or intending to take off or land" (Byrne & Kirlik, 2005; Wald, 2006). Such errors are serious enough that they are tabulated in a national-wide database and were the focus of a NASA funded effort to find a systematic explanation for their occurrence.

The modeling effort by Byrne and Kirlik had three key components. First, interviews with pilots to elicit the knowledge and strategies used to navigate from the runway to gate. Second, an analysis of 284 routes from landing to gate at 9 major U. S. airports. Third, an analysis of the time-horizon for making a decision to turn an aircraft as a function of intersection distance and aircraft dynamics (when on the ground, large aircraft are neither graceful nor easy to turn).

From their knowledge engineering effort, the researchers obtained five heuristics that pilots used to determine which direction to turn at an intersection. The cognitive modeling focused on strategy selection as a function of the time remaining before a decision had to be made and a turn initiated. As in the UAV example, an important component of the model was aircraft dynamics; however, in the passenger jet case it was the dynamics of a lumbering 870,000 lbs passenger jet as it taxied on the ground towards its gate – specifically, an algorithm "to calculate the maximum speed with which a turn of a given type could be negotiated." The model was based on the ACT-R cognitive architecture but run as a Monte Carlo simulation for 300 repetitions at each of 50 time horizons. The results predicted the selection probability for each of 5 heuristics. The details of the model and the heuristics are beyond the scope of the current paper, but a glance at Figure 2 shows large differences between predicted strategy as a function of the decision-time horizon.

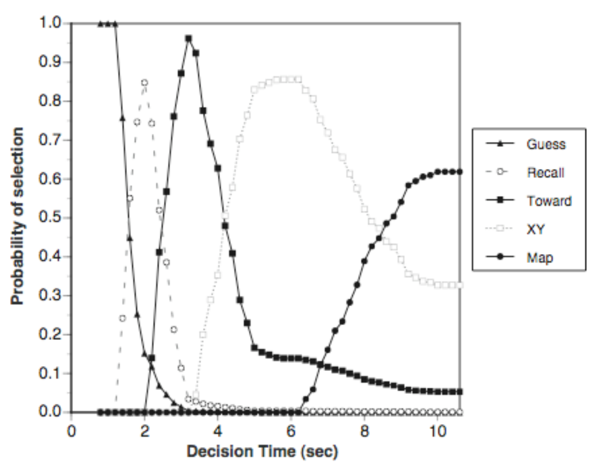

Figure 2: Predicted selection probability for each turn heuristic by decision-time horizon (from Byrne & Kirlik, 2005).

The Byrne and Kirlik work focused on errors in decision-making and showed a good match to the data set of real errors as well as to a better documented set of errors collected from experienced airline pilots in a high-fidelity flight simulator. Making an error while taxiing is a low-probability event. The chance of a serious, loss-of-life, incident following from such an error is an even lower-probability event. However, in an ATC system that supervises thousands of takeoffs and landings each day, even extremely low-probability events may come to pass. The low actual probability makes the empirical data difficult to collect as the event is occurring.

When a terrible, low-probability event does occur, all too often the public and official response is to find someone to blame – in aviation this usually means blaming the pilot or the air traffic controller. Model-based analysis such as Byrne and Kirlik's do not eliminate the human responsibility from the accident equation, but they do show that properties of the designed task environment contribute to such accidents in all-too-predictable ways. The simplest and oldest way to augment human cognition is by the design of the task environment. Byrne and Kirlik's analyses point to how this might be done to reduce errors while taxiing.

3.1.3 Driving and Driving while Dialing

An especially notable attempt to model a complex task is Salvucci's program of research to model driving (Salvucci, 2001; Salvucci & Macuga, 2001; Salvucci & Gray, 2004; Salvucci, 2006). Driving is such an everyday, or mundane, expertise that it may be necessary to step back and remind the reader of its cognitive complexity.

> The task of driving is in fact an ever-changing set of basic tasks that must be integrated and interleaved. Michon (1985) identifies three classes of task processes for driving: operational processes that involve manipulating control inputs for stable driving, tactical processes that govern safe interactions with the environment and other vehicles, and strategic processes for higher-level reasoning and planning. Driving typically involves all three types of processes working together to achieve safe, stable navigation — for instance, monitoring a traffic light and controlling the vehicle to stop and start, or deciding to make a turn and controlling the vehicle through the turn. Some tasks are not continual but intermittent, arising in specific situations — for instance, parking a vehicle at a final destination. In addition, driving may include secondary tasks, perhaps related to the primary driving task (e.g., using a navigation device), or perhaps mostly or entirely unrelated (e.g., tuning a radio or dialing a cellular phone) (Salvucci, 2006).

Salvucci's modeling work has evolved over the last several years. The most complete report of his work is contained in the *Human Factors* paper (Salvucci, 2006). In that work Salvucci presents the results of comparing the models with human behavior on several dependent variables related to lane keeping, curve negotiation, and lane changing. The dependent variables include performance-based measures such as steering angle and lateral position, as well as eye data measures thought to be closely related to visual attention. Salvucci's modeling is done within the ACT-R architecture of cognition. In all cases, his work benefits from ACT-R's ability to use the same simulation software as used by his human subjects. This means that the same type of log files are generated by models and humans, and the same types of analyses can be easily applied.

Salvucci's basic model of driving has been integrated with models of dialing different cell phones (Salvucci, 2001). The results have practical implications in yielding clear predictions for the differential effect of cell phone design on driving performance. The efforts to integrate models of two individual tasks (driving and dialing) also have implications for the control of integrated cognitive systems (Salvucci, 2005). This later work nicely illustrates the importance of cognitive engineering for identifying important gaps in basic cognitive theory.

3.1.4 Predicting Skilled Performance from Analyses of Novices

Verifying a task analysis of a complex system is itself a complex task with little guidance and few shortcuts suggested by the literature (Kirwan & Ainsworth, 1992; Shepherd, 2001). In some sense, the knowledge contained in the computational cognitive model can be considered a verification of a task analysis if the model is able to perform the task at human levels of competence. Perhaps a more stringent validation of a task analysis is that, with experience at the task, the knowledge taught to human novices suffices to produce expert level performance. Taatgen and Lee (2003) take up this challenge with modest, but noteworthy results.

The acquisition of skilled performance is an enduring topic in both the basic and applied literature. Indeed, the dominant characterization of skilled performance as passing through three phases (Fitts & Posner, 1967) comes from a researcher, Paul Fitts, who made contributions to both experimental psychology and human factors research. Taatgen and Lee embrace Anderson's (1982) characterizations of the improvement in performance between stages as a shift from declarative to procedural knowledge with practice. The model they built was written in ACT-R and was seeded with basic knowledge (how to click on buttons, where the various information items were on the display and what they represented) as well as a declarative representation of the instructions to human subjects.

The vehicle they chose for their test is the Kanfer-Ackerman air traffic controller (KA-ATC) task (Ackerman & Kanfer, 1994); a game-like simplification of the air traffic controller's task. The model used Taatgen's production compilation enhancement (Taatgen & Anderson, 2002) to ACT-R. With experience, production compilation converts declarative knowledge and very general productions into task specific productions. Through production compilation many fewer production cycles are required to perform the task and many retrievals from memory are replaced by incorporating specific declarative knowledge into the specialized productions.

The Taatgen and Lee model incorporated the knowledge and declarative representation of procedures that humans would acquire through a through reading of task instructions. All parameters used by the model were the default ACT-R parameters. The model played the game for 10, 10-min periods (trials). Model and human performance were compared at three increasingly detailed levels of analysis; overall performance, unit task performance, and keystroke levels. Across all three levels of comparison, the model mimicked the qualitative changes in human performance (as shown by generally high r^2 values). The absolute match of the model to human data (such as would be tested by RSME comparisons – though these were not provided by the authors) showed mixed results. In several cases, the models were right on top of the human data. For overall performance, the models started out better than the humans, but began to match human performance after the fifth 10-min trial. For two of the three unit tasks, the model provided a good match to human performance. In the third unit task the model was slower. For the keystroke data presented, the model also was generally slower than the humans.

The authors attribute the differences in model and human performance to several factors. First, other analyses (John & Lallement, 1997; Schunn & Reder, 2001) have established that different humans bring different strategies to bear on this task; whereas the model only used one strategy. Furthermore, human performance as it speeds up seems to exhibit more parallelism than did model performance. These differences reflect limits of the current ACT-R architecture that does not learn new strategies with experience or by exploration and can only deploy the strategies provided by the modeler. On the other hand, the general quantitative fit of model to data is good and the qualitative fit captures the speed up in human performance with experience. In general, the Taatgen and Lee work is a successful demonstration of the ability of cognitive engineering to predict expert patterns of performance given a novice task analysis and instructions.

3.1.5 Summary: Complex Systems

The above four cases are good examples of cognitive engineering applied to complex systems. Although the applications may seem modest, each of the computational cognitive models is built on an architecture, ACT-R, that was intended as a vehicle for basic research, not applied. In this usage, the ACT-R architecture becomes a vehicle for applying basic research to applied problems.

A question arises as to "why ACT-R?" The answer is straight forward and rest more on the nature of cognitive engineering than they do on ACT-R's claims as an architecture of cognition (though this aspect certainly does not hurt!). First, compared with connectionist models, it is much easier to write ACT-R models that interact with tasks that consist of multiple components that extend serially in time (e.g., such a piloting a UAV). Second, compared with Soar, over the past 15-yrs ACT-R (like EPIC) has changed dramatically to incorporate many of the theoretical advances of cognitive science. Of prime importance for modeling interactive behavior has been ACT-R's (like EPIC's) emphasis on modeling the interaction of cognitive, perceptual, and motor operations. Third and probably foremost, the current ACT-R software (ACT-R 6.0, see Anderson et al., 2004) reflects over a decade of software engineering. One result of this software engineering is that it is now relatively easy to incorporate new modules into ACT-R to either supplement or replace existing modules. Fourth, in the last 8 years (Byrne & Anderson, 1998), a significant part of the software engineering effort has been focused on enabling ACT-R models to directly interact with the same software task environments as human users. Fifth, a concomitant of this software engineering has been a decade long effort to produce tutorial materials and conduct a series of summer schools to train a wide variety of users in ACT-R.

Whatever the particular merits of ACT-R, it is clear that the last two decades have seen a significant expansion in the scope of cognitive engineering. We have gone from a focus on small, self-paced tasks such as text editing (Card et al., 1983) to larger and much more dynamic tasks such as piloting UAVs and driving cars.

3.2 Visual Analytics: Human-Computer Interaction during the Search and Exploration of Massive Sets of Qualitative Data

Visual analytics is a new label (Thomas & Cook, 2005; Wong & Thomas, 2004) for efforts to present abstract information in visual form in both structured and unstructured displays. Structured displays include the traditional bar charts and line graphs (Wainer & Velleman, 2001) as well as newer technologies that allow us to dynamically create multiple, multidimensional, complex representations of selected subsets of vast data sets (for an excellent sampling of recent innovative visualizations for visual analytic displays see the special issue organized by Keim, Robertson, Thomas, & van Wijk, 2006). The most common example of unstructured displays is the World Wide Web. Although individual web pages or web sites may be well structured, the web as a whole is not.

Motivation for work in this area is high. Knowledge may be power, but data is not knowledge until it can be processed and presented so that a human can understand and use it. Techniques for displaying data are key to transforming data into knowledge. These techniques need to support the human user in rapid, exploratory search, and in comprehending what is found. Unfortunately, modern techniques for visualization are not only prey to well known usability problems (Andre & Wickens, October 1995) but can introduce new and debilitating distortions. For example, 3-D representations of terrain, favored by many new military systems, make it extremely difficult to accurately judge relative and absolute distances (Smallman & St. John, 2005). Identifying and guarding against such alluring and subtle distortions should be one of the goals of cognitive engineering.

This section begins with a brief history of visual analytics as treated by cognitive science. Next is a sampling of recent work on structured displays. This sampling is followed by a discussion of work on seeking and extracting information from the unstructured environment of the World Wide Web.

3.2.1 A Brief History

In large part, the history of visual analytics has been driven by two very different communities; judgment and decision-making, and human-computer interaction.

3.2.1.1 Judgment and Decision-Making Beginnings

In the 80's, J. W. Payne, Bettman, and Johnson (as summarized in their 1993 book) acknowledged borrowing the construct of elementary information processes (EIPs) from Newell and Simon (1972) to quantify the cognitive effort involved in various judgment and decision-making strategies. By demonstrating that people would adopt strategies that traded off decision-making effectiveness (accuracy) for cognitive efficiency (effort), this research had a large and important influence on both the judgment and decision-making as well as the cognitive science communities (for example, Anderson, 1990; Anderson, 1991).

The EIPs construct provides a framework for thinking about the cognition involved in various decision strategies under a variety of conditions by comparing the effectiveness and efficiency of alternative strategies that use partial information against that of a decision strategy that uses complete information. However, the EIP construct is limited in that it is not embedded in a theory of the control of cognition. For example, efficiency is assessed by simply comparing the number of EIPs used by alternative strategies with the number used by the total information strategy.

The original work ignored how the information was displayed and factors associated with the cost of information extraction, manipulation, and retention. This emphasis began to change in the 1990's. Researchers in the judgment and decision-making tradition began to focus on the influence of the organization, form, and sequence of information on strategy selection (for example, Fennema & Kleinmuntz, 1995; Kleinmuntz & Schkade, 1993; Schkade & Kleinmuntz, 1994). Other research looked at how individual differences in working memory capacity interacted with interface design to affect performance on decision-making tasks (Lohse, 1997). At least one study investigated how the cost of information access affects strategy selection (Lohse & Johnson, 1996). Other studies looked at how the design of decision aids may have unintended consequences for the decision strategies that people adopt (Adelman, Bresnick, Black, Marvin, & Sak, 1996; Benbasat & Todd, 1996; Rose & Wolfe, 2000; Todd & Benbasat, 1994, 1999, 2000).

3.2.1.2 Human-Computer Interaction Beginnings

The growth of research in cognitive modeling of information search provides a case study on the role of technology in scientific research. Although other work focused on the cognitive factors implications of interface design (Gray et al., 1993; Lohse, 1993), few groups had access to the advanced technologies that were creating new designs and new visualizations (see Card, Mackinlay, & Shneiderman, 1999, for a compendium of many of the key papers from the 80's and 90's on this topic.). In the early 90's many advanced visualizations were emerging from PARC and the cognitive modeling group led by Card took full advantage of these opportunities (Card, Pirolli, & Mackinlay, 1994; Mackinlay, Robertson, & Card, 1991; Pirolli & Card, 1995; Russell, Stefik, Pirolli, & Card, 1993).

An early project focused on the "cost-of-knowledge" for extracting information from dynamic, visual analytic displays (Card et al., 1994). This effort framed the information search problem in terms of the number of documents that could be found within a given time period. Different technologies allowed different efficiencies with "technologies" widely defined to include everything from computer programs, to stacks on the desktop, to papers tucked away in filing cabinets (see Figure 3). The work identified several ways in which such searches could be sped up including alternative ways of organizing the documents, making more efficient methods available to the user, and advances in computer technologies (hardware and networks) that would allow the same procedures to be performed more quickly. These analyses illustrated the effect that tradeoffs in system design issues, very broadly defined, would have on human productivity.

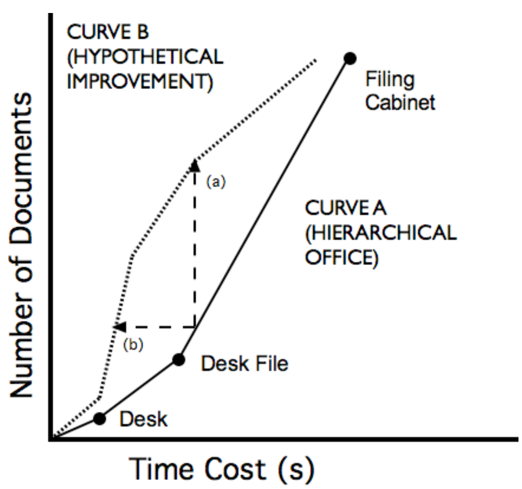

Figure 3: Cost of Knowledge Characteristic Function (adapted from Card et al., 1994, Figure 1). Notice that improvement (Curve B, the dotted line) can come in one of two ways. By keeping the time costs the same we can access more documents (arrow A). By keeping the number of documents the same we can access them at a lower time cost.

3.2.2 Reasoning from Graphs

Understanding information presented in line graphs may be the prototypical example of reasoning from structured displays. At the very least, it is a topic that has been well studied for a number of years (many of the key papers are listed in the bibliography provided by Gillan, Wickens, Hollands, & Carswell, 1998; Tufte, 1983). However, people still have trouble extracting information from graphs and, as anyone who has had to perform this task might suspect, researchers still have problems creating graphs that make their data transparent to their readers.

There are two basic decisions made in creating a data graph, (a) what to display and (b) how to display it. Most discussions of line graphs assume that the display incorporates the intended data and no more. If the "what to display" assumption is correct than alternative displays of the same data would be informationally equivalent (Larkin & Simon, 1987; Peebles & Cheng, 2003) in that no information could be inferred from one that could not be inferred from another. Of course, the assumption that the graph designer has displayed no more and no less information than needed can be problematic. Displaying too little information means that the "point" of the display can never be taken. Displaying too much means that the reader may be confused as to the point or may extract information irrelevant to the point.

If we assume that no more and no less information than is needed is being displayed then "how to display it" becomes an important topic. Alternative representations may be informationally equivalent without being computationally equivalent. Computational equivalence (Larkin & Simon, 1987) refers to the number of operations, the resources required by these operations (e.g., memory, attention, perceptual-motor), or the time required to extract information from the graph. Informational equivalence is a function of the displays whereas computational equivalence is a function of the cognitive, perceptual, and motor operations required to extract equivalent information from the displays.

Peebles and Cheng (2003) studied the basic graph reading task (see also, Lohse, 1997) of determining the value of one variable that corresponds to the given value of another variable (e.g., for the year in which oil consumption was 6, what was gas consumption?). They created two informationally equivalent versions of each graph and for each one asked three types of questions of each participant. Their intention was two-fold. First, they wished to compare human behavior in this task with the predictions of their *Graph-Based Reasoning* model (Peebles, Cheng, & Shadbolt, 1999). The Graph-Based Reasoning predictions were based on a task analysis, which assumes that the eye movements made will follow the optimal sequence required to achieve the current informational goal. Second, they wished to evaluate the value-added of building a computational cognitive model of embodied cognition that incorporated detailed assumptions regarding the use of memory, visual attention, motor m(Gray, 2006)ovement, and perceptual operations. The model was built using ACT-R. An example of the data they collected from their human subjects and their model subjects is presented as Figure 4.

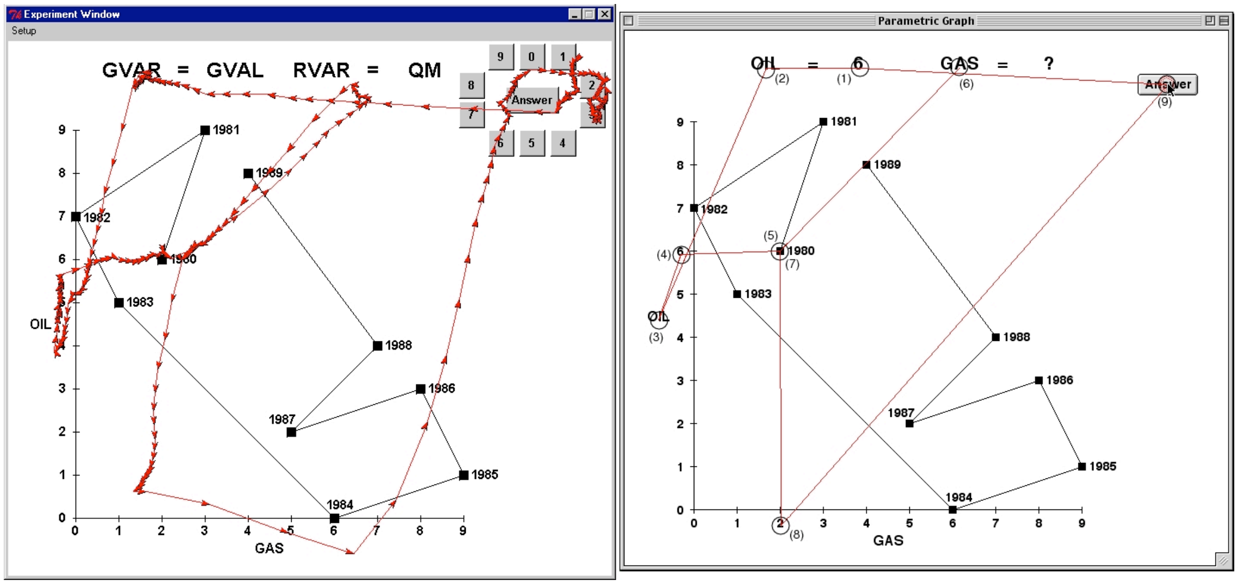

Figure 4: Screen shots from Peebles & Cheng (2003) showing a human subject's eye movement path (left) and the ACT-R models visual attention path (right). Both scan paths are in pursuit of the answer to a specific query. (For a discussion of the query and the layout of this particular graph see Peebles & Cheng, 2003.)

Peebles and Cheng found an interaction between question type and graph type. The efficiency with which questions were answered varied across the two graph types. The eye movement data revealed patterns that could be interpreted as due to perceptual and cognitive limitations of the participants. The Graph-Based Reasoning model did not predict these patterns; however, the ACT-R model did. Compared to predictions from the Graph-Based reasoning model, the ACT-R model required more glances back to the question toward the end of each trial. These glances varied with question type and graph type and reflect different demands made on memory for the location of graph elements, the association between symbol and variable, and the effect of practice.

Traditional approaches to graph understanding have focused on two elements: the visual properties of the graph and the requirements of the task. The model built by Peebles and Cheng incorporates a third element; namely, the demands on embodied cognition (cognitive, perceptual, and motor operations). The interaction among these three elements is complex. Compared to the other tasks discussed in this chapter, the familiarity of graph reading to the readers of this chapter may make it seem a cut and dried proposition. However mundane the task may be, it is clearly a task that even the most experienced of scientists struggles with when reading the work in their field. It may not be too far-fetched to think that the day may be near when each author could have a graph-reading program that would provide real-time feedback as to how much difficulty a novice (one unfamiliar with the terminology) and expert would have in extracting different amounts of information from a display. If this comes to pass it would be through the application of cognitive science theory for cognitive engineering purposes.

3.2.3 Cognitive Engineering Models of Surfing the Web (and other informational displays): Information Search of Massive Data Sets in Heterogeneously Designed Spaces

A prime example of unstructured displays for representing and accessing massive amounts of quantitative data visually is the World Wide Web. Although surfing the web represents a mundane expertise (i.e., something that many people in the population of readers do daily and are very good at), understanding and predicting human web search behavior presents a considerable challenge to the theoretical mechanisms of contemporary cognitive science. In particular, a full accounting of web search requires an understanding of information scent, semantic relatedness, visual saliency, general knowledge of page layouts and idioms, as well as the control of integrated cognitive systems.

3.2.3.1 Theory Development

Analogies are important in making progress in science. In the mid-90's, Pirolli hit on the analogy between people searching for information and animals searching for food. This tie between information search and animal foraging was productive due, in large part, to the existence of *optimal foraging theory* (Stephens & Krebs, 1986) that cast animal foraging behavior into an abstract, quantitative, modeling framework that proved possible to adapt and extend to human information seeking (Pirolli & Card, 1999).

By this analogy, *information scent* is the construct that relates the semantic relatedness of the current information context to the goals of the user's search. The higher the information scent, the more related the current context is to the search goals and the more likely the searcher would be to remain in the current *information patch*. The searcher will likely remain at the current patch for as long as the information scent remains above some threshold. When the information scent falls below this threshold the searcher will likely leave (Card et al., 2001; Fu & Pirolli, 2007; Pirolli & Fu, 2003). In web terms, a query to a search engine may return a link with a high information scent to the user's query. The user will click on the link and begin to search the web site. As long as the information at that site has a high semantic relatedness to the searched-for information, the user will remain. However, if the information turns out to be less useful than expected, the user will leave the current site and attempt to find one with a higher information scent.

The theory of information scent has played a key role in developing the cognitive engineering approach to information search and retrieval. However, it is the first step, not the last in building cognitive engineering models of human information search. If our goal is to design cognitively congruent interactive procedures for searching information displays, then nothing less than a *cognitive theory of embodied information search* is required. Some of the issues and cognitive technologies for achieving this goal are discussed in the following sections.

3.2.3.2 Semantic Relatedness Measures

For the development of models of information search, an all-but-prerequisite co-development was that of statistical measures of semantic relatedness (Harman, 1993; Landauer & Dumais, 1997; Lemaire & Denhiére, 2004; Turney, 2001). To appreciate the development of these measures, consider how the modeling of interactions with semantic content would have been handled without these measures.

For decades, the only means available within experimental psychology of estimating associative strength was to have participants in psychology studies estimate the link between two items on, say, a seven point scale. By this method, obtaining reliable and valid estimates of associative strength between, say, each of 100 words would require human judgments on 4,950 word pairs. As such human judgments are notoriously noisy, a reliable and valid estimate of relatedness required large numbers of human subjects to judge the associative relatedness of each pair of words. At the end of all of this work, there would be good estimates of relatedness between each of the 100 words in the list. Obviously, such methods of obtaining word associations means that human search among the unlimited diversity of the WWW would be all but impossible to model and study.

Statistical measures of semantic similarity parse large corpora (measured in the millions of documents) to develop families of measures that purport to provide human-like judgments of the relatedness of one word to another or of an entire text to the topic of a query. A review of these measures is beyond the scope of the current chapter. (See, Lemaire & Denhiére, 2004, for a cogent and succinct overview of several of these measures.) However, it is clear that the research in this area has gone from demonstrating that various measures can be generated that mimic human judgment, to examining the cognitive fidelity of these statistically derived judgments across a variety of search tasks (Blackmon et al., 2005; Kaur & Hornof, 2005; Veksler & Gray, 2006). The current goals include understanding the limits and best application of current methods (Juvina, van Oostendorp, Karbor, & Pauw, 2005; Lemaire & Denhiére, 2004; Matveeva, Levow, Farahat, & Royer, 2005) and developing new measures as needed.

3.2.3.3 Stimulus-Driven Factors in Visual Search

The human visual system seems hard-wired by evolution to allocate attention and resources to certain visual stimuli (Wolfe, 1998) such as abrupt onsets or patterns of motion against otherwise still backgrounds (Ware & Bobrow, 2004). It also seems the case that certain combinations of stimulus characteristics are easier to allocate visual attention to in some visual environments than in others. For example, searching for a red L amid a field of green L's is quite easy. Searching for a red L amid a field of green L's and red T's is much harder (Wolfe, 1998).

Recent work has made progress in developing statistical visual saliency measures (Itti & Koch, 2001; Rao & Ballard, 1995; Rosenholtz, 2001; Verghese & McKee, 2004) that are somewhat analogous to the statistical semantic relatedness measures discussed above. These measures allow the computation of differences between elements of a visual display along one or more visual dimensions. For example, Figure 5 shows a visual saliency map of a webpage using Rosenholtz' (Rosenholtz, 1999) measure of visual saliency. This measure computes a similarity score for each of three dimensions (color, orientation, and contrast) between every element on the screen and every other element.

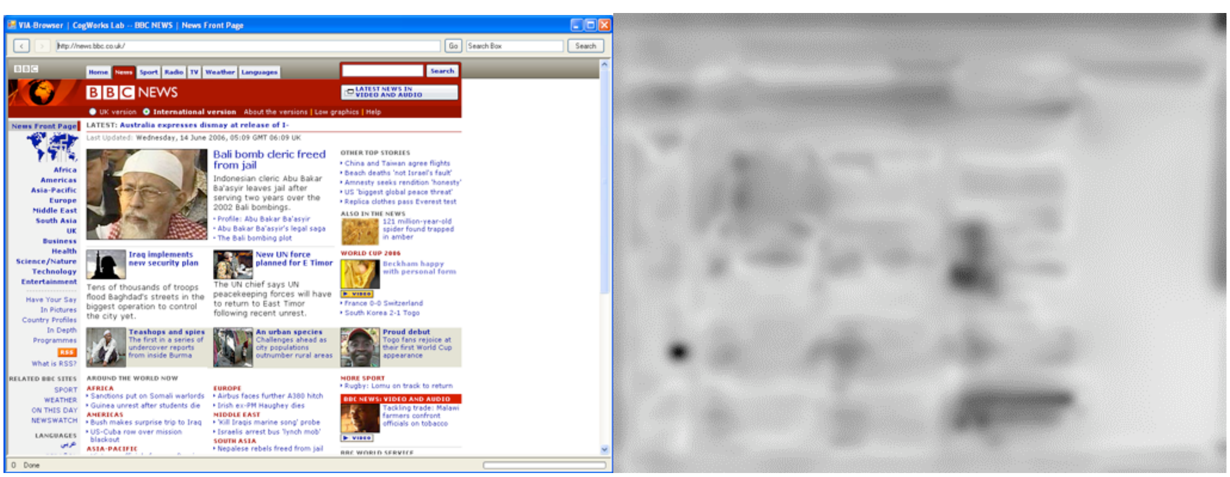

Figure 5: Visual Saliency Map on right of webpage shown on left using Rosenholtz' (1999) measure of visual saliency. To produce the "map", for each screen pixel, the score for each dimension is translated to a 256-bit vector, averaged with the other two dimensions and mapped onto a 256 bit gray-scale.

It is certainly not the case that visual saliency determines where the eye looks. If this issue were in doubt, it has been laid to rest by a variety of small, experimental psychology demonstrations (Henderson, Brockmole, Castelhano, & Mack, in press; Underwood & Foulsham, 2006). However, contemporary theory asserts that eye movement location is determined by activation at a saccade map (Trappenberg, Dorris, Munoz, & Klein, 2001) and that where the activation builds up is influenced by stimulus-driven as well as goal-directed factors (Godijn & Theeuwes, 2002; Pomplun, 2007). It seems clear that visual saliency has an important role to play in the direction, or misdirection, of visual attention during information search. Given the visually dramatic nature of modern structured and unstructured (e.g., the left-side of Figure 5**Error! Reference source not found.**) visual analytic displays, it is clear that visual saliency is a factor that must be taken into account in predicting the success and time costs of information search.

3.2.3.4 Knowledge-Driven Factors in Visual Search

Knowledge-driven factors are those that influence visual search because of explicitly *adopted* strategies (Everling & Fischer, 1998; Godijn & Theeuwes, 2002), *adapting* search to the statistical structure of a specific search environment (Reder, Weber, Shang, & Vanyukov, 2003), or general knowledge and experience with the layout of visually presented information (e.g., books have tables of content, indexes, and chapters; chapters have introductions, subsections, and summaries; and so on). It is clear that even unstructured information spaces such as web pages often enable knowledge-driven search. For example, most web pages have headers and footers that contain navigation information or information specific to the page or website itself; that is, not content information. Likewise, menus come in many forms, but most of these forms are visually distinctive and, when present, are often on the right, the left, or towards the top of a webpage (Rigutti & Gerbino, 2004). At present, we know of no computational cognitive model that factors such knowledge into its predictions, though there are several in which such factors are explicitly eliminated by the modelers (Blackmon et al., 2005; Kaur & Hornof, 2005).

3.2.4 Summary: Cognitive Engineering Models for Human-Information Interaction

The purpose of building cognitive engineering models that can surf the web, extract information from line graphs, or make sense of a complex visual display of qualitative information is to optimize human performance by identifying design elements or decisions that may lead to stable but suboptimal human performance (Fu & Gray, 2004, 2006; Fu, 2007). Achieving this goal will require advances in cognitive science theories of semantic comprehension and visual attention, as well as in advances in our understanding of the composition and control of integrated cognitive systems (Gray, 2007a). The symbiosis between cognitive science and cognitive engineering shows no sign of abating.

4 Conclusions: Themes and Challenges

This chapter began with a discussion of five dimensions on which cognitive science and cognitive engineering sometimes differ.

First discussed was the manner in which the problems were picked. As in the case of runway incursions (Byrne & Kirlik, 2005; Kirlik, 2007) or UAVs (Gluck et al., 2007), the problem is often picked for the researcher and the research is unlikely to continue without a larger organization that provides substantive support in terms of equipment and specialized expertise.

Second was the amount of prior empirical data. Certainly the models discussed above rest on a broad base of empirical data. However, the breadth of that base meant that much of it, usually the part closest to the task being modeled, was very shallow. Hence the Peebles and Cheng model (Peebles & Cheng, 2003) rested on years of research on memory and visual attention as well as studies of reading line graphs. However, theirs is the first model that put these elements together to predict graph reading. The shallowness of this base was fully exposed in the discussion of Visual Analytics in which a common everyday task (i.e., surfing the web) presents a challenge to both applied and basic cognitive science.

Third was the expertise factor. Key to modeling the UAV, runway incursions, and driving tasks was access to human experts. This factor was not as important to the graph-reading task and not at all applicable to the ATC game where the focus was on understanding the transition from novice to skilled performance. It is unclear how this dimension of expertise will play itself out in the domain of Visual Analytics. Certainly the strong reliance on statistical measures of semantic relatedness suggest that much mundane human expertise will have to be assumed before progress on building cognitive engineering models of Visual Analytic performance can progress.

Fourth was a different sort of factor. The areas in which cognitive engineering has been applied are areas with strong demands for answers. What are the implications for cognitive science if cognitive engineering cannot meet those demands? If cognitive engineering is “cognitive science theories applied to human factors practice”, than if we cannot meet the demands for cognitive engineering, a strong implication might be that the cognitive science enterprise has little relevance to our modern world. Pursuit of knowledge for knowledge’s sake may be a fine ideal, but the successful advancement of physics and biology in the last century was due less to knowledge for knowledge’s sake and more to knowledge that was able to address key concerns of our society. Compared to disciplines such as Philosophy and Linguistics, research in cognitive science is well funded. The basic questions that Cognitive Science addresses are no more compelling than those in these other fields. What accounts for our better funding is our promise to society that advances in cognitive science theory will result in tangible improvements.

Fifth was the view that cognitive engineering dealt with integrated cognitive systems and that such systems had been largely ignored by a basic research community that was content to dive deeply into artificially isolated areas such as reasoning, decision making, memory, and visual attention (Byrne, 2007b). Not only is the control of integrated cognitive systems a challenging basic research question, the importance of understanding the control of integrated cognitive systems for cognitive engineering purposes suggest that research on control issues should become a high priority among basic researchers as well as those agencies that fund basic research.

References

Ackerman, P. L., & Kanfer, R. (1994). Kanfer-Ackerman air traffic controller task© CD-ROM database, data collection program, and playback program. Arlington, VA: Office of Naval Research, Cognitive Science Program.

Adelman, L., Bresnick, T., Black, P. K., Marvin, F. F., & Sak, S. G. (1996). Research with Patriot Air Defense officers: Examining information order effects. *Human Factors, 38*(2), 250–261.

Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review, 89*(4), 369-406.

Anderson, J. R., Conrad, F. G., & Corbett, A. T. (1989). Skill acquisition and the LISP Tutor. *Cognitive Science, 13*(4), 467–505.

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences, 14*(3), 471–517.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2*, 396–408.

Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J. R., Matessa, M., & Lebiere, C. (1997). ACT-R: A theory of higher-level cognition and its relation to visual attention. *Human-Computer Interaction, 12*(4), 439–462.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglas, S., Lebiere, C., & Quin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111*(4), 1036-1060.

Andre, A. D., & Wickens, C. D. (October 1995). When users want what's not best for them. *Ergonomics in design,* 10–14.

Ballard, D. H., Hayhoe, M. M., & Pelz, J. B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience, 7*(1), 66–80.

Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences, 20*(4), 723–742.

Ballard, D. H., & Sprague, N. (2007). On the role of embodiment in modeling natural behaviors. In W. D. Gray (Ed.), *Integrated models of cognitive systems*. New York: Oxford University Press.

Benbasat, I., & Todd, P. (1996). The effects of decision support and task contingencies on model formulation: A cognitive perspective. *Decision Support Systems, 17*(4), 241–252.

Blackmon, M. H., Kitajima, M., & Polson, P. G. (2005). Tool for accurately predicting website navigation problems, non-problems, problem severity, and effectiveness of repairs. In W. Kellogg, S. Zhai, G. v. d. Veer & C. Gale (Eds.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 31-40). New York: ACM Press.

Bovair, S., Kieras, D. E., & Polson, P. G. (1990). The acquisition and performance of text-editing skill: A cognitive complexity analysis. *Human-Computer Interaction, 5*(1), 1–48.

Byrne, M. D., & Anderson, J. R. (1998). Perception and action. In J. R. Anderson & C. Lebiére (Eds.), *The atomic components of thought* (pp. 167–200). Hillsdale, NJ: Lawrence Erlbaum Associates.

Byrne, M. D., & Kirlik, A. (2005). Using computational cognitive modeling to diagnose possible sources of aviation error. *International Journal of Aviation Psychology, 15*, 135-155.

Byrne, M. D. (2007a). Cognitive architecture. In A. Sears & J. A. Jacko (Eds.), *The human-computer interaction handbook* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.

Byrne, M. D. (2007b). Local theories vs. comprehensive architectures: The cognitive science jigsaw puzzle. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 431-443). New York: Oxford University Press.

Card, S. K., Moran, T. P., & Newell, A. (1980a). Computer text editing: An information processing analysis of a routine cognitive skill. *Cognitive Psychology, 12*, 32–74.

Card, S. K., Moran, T. P., & Newell, A. (1980b). The keystroke-level model for user performance time with interactive systems. *Communications of the ACM, 23*(7), 396–410.

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

Card, S. K., Pirolli, P., & Mackinlay, J. D. (1994). The cost-of-knowledge characteristic function: Display evaluation for direct-walk dynamic information visualizations. In B. Adelson, S. Dumais & J. Olson (Eds.), *ACM CHI'94 Conference on Human Factors in Computing Systems* (Vol. 1, pp. 238–244). New York: ACM Press.

Card, S. K., Pirolli, P., Van Der Wege, M., Morrison, J. B., Reeder, R. W., Schraedley, P. K., et al. (2001). Information scent as a driver of Web behavior graphs: results of a protocol analysis method for Web usability. *CHI Letters, 3*(1), 498–505.

Card, S. K., Mackinlay, J. D., & Shneiderman, B. (Eds.). (1999). *Information visualization: Using vision to think*. New York: Morgan Kaufmann Publishers.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*(2), 121–152.

Chipman, S. F., & Kieras, D. E. (2004, March 17-18). *Operator centered design of ship systems.* Paper presented at the Engineering the Total Ship Symposium, 2004, NIST, Gaithersburg, MD.

Corbett, A. T., & Anderson, J. R. (1988). The LISP intelligent tutoring system: Research in skill acquisition. In J. Larkin, R. Chabay & C. Scheftic (Eds.), *Computer assisted instruction and intelligent tutoring systems: Establishing communciations and collaboration*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Ehret, B. D., Gray, W. D., & Kirschenbaum, S. S. (2000). Contending with complexity: Developing and using a scaled world in applied cognitive research. *Human Factors, 42*(1), 8–23.

Elkind, J. I., Card, S. K., Hochberg, J., & Huey, B. M. (1989). *Human performance models for computer-aided engineering*. Washington, DC: National Academy Press.

Everling, S., & Fischer, B. (1998). The antisaccade: a review of basic research and clinical studies. *Neuropsychologia, 36*(9), 885-899.

Fennema, M. G., & Kleinmuntz, D. N. (1995). Anticipations of effort and accuracy in multiattribute choice. *Organizational Behavior & Human Decision Processes, 63*(1), 21–32.

Fitts, P. M., & Posner, M. I. (1967). *Human performance*. Belmont, CA: Brooks Cole.

Freed, M., Matessa, M., Remington, R. W., & Vera, A. (2003, March). *How Apex automates CPM-GOMS.* Paper presented at the Fifth International Conference on Cognitive Modeling, Bamberg, GE.

Fu, W.-T., & Gray, W. D. (2004). Resolving the paradox of the active user: Stable suboptimal performance in interactive tasks. *Cognitive Science, 28*(6), 901-935.

Fu, W.-T., & Gray, W. D. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology, 52*(3), 195-242.

Fu, W.-T. (2007). A rational–ecological approach to the exploration/exploitation trade-offs: Bounded rationality and suboptimal performance. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 165-179). New York: Oxford University Press.

Fu, W.-T., & Pirolli, P. (2007). SNIF-ACT: A cognitive model of user navigation on the World Wide Web. *Human-Computer Interaction, in press*.

Gillan, D. J., Wickens, C. D., Hollands, J. G., & Carswell, C. M. (1998). Guidelines for presenting quantitative data in HFES publications. *Human Factors, 40*(1), 28–41.

Gluck, K. A., Ball, J. T., Gunzelmann, G., Krusmark, M. A., Lyon, D. R., & Cooke, N. J. (2005). *A Prospective Look at Synthetic Teammate for UAV Applications.* Paper presented at the American Institute of Aeronautics and Astronautics "Infotech@Aerospace" Conference, Reston, VA.

Gluck, K. A., & Pew, R. W. (2005). *Modeling human behavior with integrated cognitive architectures: Comparison, evaluation, and validation* Mahwah, NJ: Lawrence Erlbaum Associates.

Gluck, K. A., Ball, J. T., & Krusmark, M. A. (2007). Cognitive control in a computational model of the predator pilot. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 13-28). New York: Oxford University Press.

Godijn, R., & Theeuwes, J. (2002). Programming of endogenous and exogenous saccades: Evidence for a competitive integration model. *Journal of Experimental Psychology: Human Perception and Performance, 28*(5), 1039-1054.

Gray, W. D., John, B. E., & Atwood, M. E. (1993). Project Ernestine: Validating a GOMS analysis for predicting and explaining real-world performance. *Human-Computer Interaction, 8*(3), 237–309.

Gray, W. D., Kirschenbaum, S. S., & Ehret, B. D. (1997). The précis of Project Nemo, phase 1: Subgoaling and subschemas for submariners. In M. G. Shafto & P. Langley (Eds.), *Nineteenth Annual Conference of the Cognitive Science Society* (pp. 283–288). Hillsdale, NJ: Lawrence Erlbaum Associates.

Gray, W. D. (2000). The nature and processing of errors in interactive behavior. *Cognitive Science, 24*(2), 205–248.

Gray, W. D., & Boehm-Davis, D. A. (2000). Milliseconds Matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experimental Psychology: Applied, 6*(4), 322–335.

Gray, W. D., & Kirschenbaum, S. S. (2000). Analyzing a novel expertise: An unmarked road. In J. M. C. Schraagen, S. F. Chipman & V. L. Shalin (Eds.), *Cognitive task analysis* (pp. 275–290). Mahwah, NJ: Lawrence Erlbaum Associates.

Gray, W. D., & Pew, R. W. (2004). Introduction to human performance modeling (HPM) and to this symposium on cognitive HPM. In *48th Annual Conference of the Human Factors and Ergonomics Society* (pp. 2109-2110). Santa Monica, CA: Human Factors and Ergonomics Society.

Gray, W. D., Schoelles, M. J., & Myers, C. W. (2004). Meeting Newell's other challenge: Cognitive architectures as the basis for cognitive engineering. *Behavioral & Brain Sciences, 26*(5), 609-610.

Gray, W. D. (2006, April 20). *The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior.* Paper presented at the Air Force Office of Scientific Research: Cognitive Program Review, Dayton, OH.

Gray, W. D. (2007a). Composition and control of integrated cognitive systems. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 3-12). New York: Oxford University Press.

Gray, W. D. (Ed.). (2007b). *Integrated models of cognitive systems*. New York: Oxford University Press.

Green, P. A. (1999). Estimating compliance with the 15-second rule for driver-interface usabilty and safety. In *Proceedings of the Human Factors and Ergonomics Society 43rd Annual Meeting.:* . Santa Monica, CA: Human Factors and Ergonomics Society.

Green, P. A. (2002). *Calculation of the time to complete in-vehicle navigation and route guidance tasks* (No. SAE J2365). Warrendale, PA: SAE International.

Harman, D. (1993). Overview of the first TREC conference, *Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval*. Pittsburgh, Pennsylvania, United States: ACM Press.

Henderson, J. M., Brockmole, J. R., Castelhano, M. S., & Mack, M. (in press). Visiual Salliency Does Not Account for Eye Movements during Visual Search in Real-World Scenes. In R. van Gompel, M. Fischer, W. Murray & R. Hill (Eds.), *Eye Movement Research: Insights into Mind and Brain*: Elsevier.

Hintzman, D. L. (2005). Memory strength and recency judgments. *Psychonomic Bulletin & Review, 12*(5), 858-864.

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience, 2*(3), 194-203.

John, B. E. (1988). Contributions to engineering models of human-computer interaction, *Unpublished doctoral dissertation*. Pittsburgh, PA: Carnegie Mellon University.

John, B. E. (1993). *A quantitative model of expert transcription typing* (No. CMU-CS-93-120): Carnegie Mellon University School of Computer Science.

John, B. E. (1996). TYPIST: A theory of performance in skilled typing. *Human-Computer Interaction, 11*(4), 321–355.

John, B. E., & Lallement, Y. (1997). Strategy use while learning to perform the Kanfer-Ackerman Air Traffic Controller task. In M. G. Shafto & P. Langley (Eds.), *Nineteenth Annual Conference of the Cognitive Science Society* (pp. 337–342). Hillsdale, NJ: Lawrence Erlbaum Associates.

Johnson-Laird, P. N. (1993). *Human and machine thinking*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Juvina, I., van Oostendorp, H., Karbor, P., & Pauw, B. (2005). Towards modeling contextual information in web navigation. In B. G. Bara, L. Barsalou & M. Bucciarelli (Eds.), *27th Annual Meeting of the Cognitive Science Society, CogSci2005* (pp. 1078-1083). Austin, Tx: The Cognitive Science Society, Inc.

Kaur, I., & Hornof, A. J. (2005). A comparison of LSA, WordNet, and PMI-IR for predicting user click behavior. In W. Kellogg, S. Zhai, G. v. d. Veer & C. Gale (Eds.), *ACM CHI 2005 Conference on Human Factors in Computing Systems* (pp. 51-60). New York: ACM Press.

Keim, D. A., Robertson, G. G., Thomas, J. J., & van Wijk, J. J. (2006). Guest editorial: Special section on visual analytics. *IEEE Transactions on Visualization and Computer Graphics, 12*(6), 1361-1362.

Kieras, D. E., & Bovair, S. (1986). The acquisition of procedures from text: A production-system analysis of transfer of training. *Journal of Memory and Language, 25*, 507–524.

Kieras, D. E. (1997). A guide to GOMS model usability evaluation using NGOMSL. In M. Helander, T. K. Landauer & P. Prabhu (Eds.), *Handbook of Human-Computer Interaction* (Second ed., pp. 733–766). New York: Elsevier.

Kieras, D. E., & Meyer, D. E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction, 12*(4), 391–438.

Kieras, D. E. (2007). Model-based evaluation. In A. Sears & J. A. Jacko (Eds.), *The human-computer interaction handbook* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.

Kirlik, A. (2007). Ecological resources for modeling interactive cognition and behavior. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 194-210). New York: Oxford University Press.

Author. (Year of Conference, Date). Title. Poster presented at the Conference Name, Conference Location.

Kirwan, B., & Ainsworth, L. K. (Eds.). (1992). *A guide to task analysis*. Washington, DC: Taylor & Francis.

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

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104*(2), 211-240.

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

Lemaire, B., & Denhiére, G. (2004). Incremental construction of an associative network from a corpus. In K. D. Forbus, D. Gentner & T. Regier (Eds.), *26th Annual Meeting of the Cognitive Science Society, CogSci2004* (pp. 825-830). Hillsdale, NJ: Lawrence Erlbaum Publisher.

Lohse, G. L. (1993). A cognitive model for understanding graphical perception. *Human-Computer Interaction, 8*(4), 353–388.

Lohse, G. L., & Johnson, E. J. (1996). A comparison of two process tracing methods for choice tasks. *Organizational Behavior and Human Decision Processes, 68*(1), 28–43.

Lohse, G. L. (1997). The role of working memory on graphical information processing. *Behaviour & Information Technology, 16*(6), 297–308.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review, 111*(2), 309-332.

Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive Psychology, 31*, 168–217.

Mackinlay, J. D., Robertson, G. G., & Card, S. K. (1991). The perspective wall: detail and context smoothly integrated. In S. P. Robertson, G. M. Olson & J. S. Olson (Eds.), *Proceedings of the SIGCHI conference on Human factors in computing systems: Reaching through technology* (pp. 173-176). New York: ACM Press.

Matveeva, I., Levow, G., Farahat, A., & Royer, C. (2005, September). *Term representation with generalized latent semantic analysis.* Paper presented at the 2005 Conference on Recent Advances in Natural Language Processing, Borovets, Bulgaria.

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

Newell, A., & Card, S. K. (1985). The prospects for psychological science in human-computer interaction. *Human-Computer Interaction, 1*(3), 209–242.

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

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 104*(2), 266-300.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.

Author. (Year of Conference, Date). Title. Poster presented at the Conference Name, Conference Location.

Peebles, D., & Cheng, P. C. H. (2003). Modeling the effect of task and graphical representation on response latency in a graph reading task. *Human Factors, 45*(1), 28–46.

Pew, R. W. (2007). Some history of human performance modeling. In W. D. Gray (Ed.), *Integrated models of cognitive systems*. New York: Oxford University Press.

Pew, R. W., & Mavor, A. S. (Eds.). (1998). *Modeling human and organizational behavior: Application to military simulations*. Washington, DC: National Academy Press.

Pirolli, P., & Card, S. K. (1995). Information foraging in information access environments. In I. R. Katz, R. Mack, L. Marks, M. B. Rosson & J. Nielsen (Eds.), *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 51-58). New York: ACM Press.

Pirolli, P., & Card, S. K. (1999). Information foraging. *Psychological Review, 106*(4), 643–675.

Pirolli, P., & Fu, W.-T. (2003). SNIF-ACT: A model of information foraging on the World Wide Web. *Lecture Notes in Computer Science, 2702*, 45-54.

Pirolli, P. (2005). Rational analyses of information foraging on the Web. *Cognitive Science, 29*(3), 343-373.

Pomplun, M. (2007). Advancing Area Activation towards a General Model of Eye Movements in Visual Search. In W. D. Gray (Ed.), *Integrated models of cognitive systems* (pp. 120-131). New York: Oxford University Press.

Rao, R. P. N., & Ballard, D. H. (1995). An active vision architecture based on iconic representations. *Artificial Intelligence, 78*(1-2), 461-505.

Reder, L. M., Weber, K., Shang, J., & Vanyukov, P. M. (2003). The adaptive character of the attentional system: Statistical sensitivity in a target localization task. *Journal of Experimental Psychology-Human Perception and Performance, 29*(3), 631-649.

Rigutti, S., & Gerbino, W. (2004). Navigating within a web site: the WebStep model. In M. C. Lovett, C. D. Schunn, C. Lebiere & P. Munro (Eds.), *Abstract presented at the Sixth International Conference on Cognitive Modeling, ICCM2004* (pp. 378-379). Mahwah, NJ: Erlbaum.

Rips, L. J. (1994). *The psychology of proof: Deductive reasoning in human thinking*: The MIT Press.

Rose, J. M., & Wolfe, C. J. (2000). The effects of system design alternatives on the acquisition of tax knowledge from a computerized tax decision aid. *Accounting Organizations and Society, 25*(3), 285–306.

Rosenholtz, R. (1999). A simple saliency model predicts a number of motion popout phenomena. *Vision Research, 39*, 3157-3163.

Rosenholtz, R. (2001). Search asymmetries? What search asymmetries? *Perception & Psychophysics, 63*(3), 476-489.

Russell, D. M., Stefik, M. J., Pirolli, P., & Card, S. K. (1993). The Cost Structure of Sensemaking. In S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel & T. White (Eds.), *ACM INTERCHI'93 Conference on Human Factors in Computing Systems* (pp. 269–276). New York: ACM Press.

Salvucci, D. D. (2001). Predicting the effects of in-car interface use on driver performance: an integrated model approach. *International Journal of Human-Computer Studies, 55*(1), 85–107.

Salvucci, D. D., & Macuga, K. L. (2001). Predicting the effects of cell-phone dialing on driver performance. In E. M. Altmann, A. Cleeremans, C. D. Schunn & W. D. Gray (Eds.), *Fourth International Conference on Cognitive Modeling* (pp. 25–30). Mahwah, NJ: Lawrence Erlbaum Associates.

Salvucci, D. D., & Gray, R. (2004). A two-point visual control model of steering. *Perception, 33*(10), 1233-1248.

Salvucci, D. D. (2005). A multitasking general executive for compound continuous tasks. *Cognitive Science, 29*(3), 457-492.

Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors, 48*(2), 362-380.

Schkade, D. A., & Kleinmuntz, D. N. (1994). Information displays and choice processes: Differential effects of organization, form, and sequence. *Organizational Behavior & Human Decision Processes, 57*(3), 319–337.

Schunn, C. D., & Reder, L. M. (2001). Another source of individual differences: Strategy adaptivity to changing rates of success. *Journal of Experimental Psychology-General, 130*(1), 59–76.

Schweickert, R., Fisher, D. L., & Proctor, R. W. (2003). Steps toward building mathematical and computer models from cognitive task analyses. *Human Factors, 45*(1), 77–103.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs, 75*(13), 42-42.

Shepherd, A. (1998). HTA as a framework for task analysis. *Ergonomics, 41*(11), 1537–1552.

Shepherd, A. (2001). *Hierarchical task analysis*. New York: Taylor & Francis.

Simon, H. A. (1996). *The sciences of the artificial* (Third ed.). Cambridge, MA: The MIT Press.

Singley, M. K., & Anderson, J. R. (1989). *The transfer of cognitive skill*. Cambridge, MA: Harvard University Press.

Smallman, H. S., & St. John, M. (2005). Naîve realism: Misplaced faith in the utility of realistic displays. *Ergonomics in Design, 13*(3), 6-13.

Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.

Taatgen, N. A., & Anderson, J. R. (2002). Why do children learn to say "broke'? A model of learning the past without feedback. *Cognition, 86*(2), 123–155.

Taatgen, N. A., & Lee, F. J. (2003). Production composition: A simple mechanism to model complex skill acquisition. *Human Factors, 45*(1), 61–76.

Thomas, J. J., & Cook, K. A. (Eds.). (2005). *Illuminating the path: The research and development agenda for visual analytics*. Washington, DC: IEEE Press.

Todd, P., & Benbasat, I. (1994). The Influence of Decision Aids on Choice Strategies - an Experimental-Analysis of the Role of Cognitive Effort. *Organizational Behavior and Human Decision Processes, 60*(1), 36–74.

Todd, P., & Benbasat, I. (1999). Evaluating the impact of DSS, cognitive effort, and incentives on strategy selection. *Information Systems Research, 10*(4), 356–374.

Todd, P., & Benbasat, I. (2000). Inducing compensatory information processing through decision aids that facilitate effort reduction: An experimental assessment. *Journal of Behavioral Decision Making, 13*(1), 91–106.

Trappenberg, T. P., Dorris, M. C., Munoz, D. P., & Klein, R. M. (2001). A model of saccade initiation based on the competitive integration of exogenous and endogenous signals in the superior colliculus. *Journal of Cognitive Neuroscience, 13*(2), 256-271.

Tsapogas, J. (2003). *Characteristics of doctoral scientists and engineers in the United States* (No. NSF 06-320). Arlington, VA: National Science Foundation.

Tufte, E. R. (1983). *The visual display of quantitative information*. Chesire, Connecticut: Graphics Press.

Turney, P. (2001). Mining the Web for Synonyms: PMI-IR versus LSA on TOEFL. In L. De Raedt & P. Flach (Eds.), *Proceedings of the Twelfth European Conference on Machine Learning (ECML-2001)* (pp. 491-502). New York: Springer.

Underwood, G., & Foulsham, T. (2006). Visual saliency and semantic incongruency influence eye movements when inspecting pictures. *Quarterly Journal of Experimental Psychology, 59*(11), 1931-1949.

Veksler, V. D., & Gray, W. D. (2006). Test case selection for evaluating measures of semantic distance. In R. Sun (Ed.), *Proceedings of the 28th Annual Meeting of the Cognitive Science Society* (pp. 2624). Austin, TX: Cognitive Science Society.

Verghese, P., & McKee, S. P. (2004). Visual search in clutter. *Vision Research, 44*(12), 1217-1225.

Wainer, H., & Velleman, P. F. (2001). Statistical graphics: Mapping the pathways of science. *Annual Review of Psychology, 52*, 305-335.

Wald, M. L. (2006, November 3). F.A.A. finds more errors on runways. *New York Times*.

Ware, C., & Bobrow, R. (2004). Motion to support rapid iteractive queries on node-link diagrams. *ACM Transactions on Applied Perception, 1*(1), 3-18.

Williams, K. E. (2000). An automated aid for modeling human-computer interaction. In J. M. Schraagen, S. F. Chipman & V. L. Shalin (Eds.), *Cognitive task analysis* (pp. 165–180). Mahwah, NJ: Lawrence Erlbaum Associates.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review, 9*(4), 625–636.

Wolfe, J. M. (1998). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13–73). East Sussex, UK: Psychology Press.

Wong, P. C., & Thomas, J. (2004). Visual analytics. *IEEE Computer Graphics and Applications, 24*(5), 20-21.

Appendix A

Different approaches to the topic of cognitive engineering and cognitive engineering models are possible and two excellent chapters that are very different from the current chapter have been authored by Kieras (2007) and Byrne (2007a).