

Using Models of Cognition in HRI Evaluation and Design*

Michael A. Goodrich

Computer Science Department
Brigham Young University
mike@cs.byu.edu

Abstract

In designing and evaluating human-machine systems, cognitive models can be used to (a) provide design principles and (b) guide the construction of experiments. In this paper, we present an information processing model of cognition that we have used extensively in designing and evaluating interfaces and autonomy modes. This model uses a conventional description of short-term memory, but treats long-term memory as a collection of mental models specific for particular tasks. Working memory includes components of both short term and long term memory; short term memory acts as a “scratch” pad for an activated subset of long-term memory. We review this model and discuss how it has been used in several human-robot systems.

I have a map of the United States . . . actual size. —
Stephen Wright, comedian.

Introduction

A model of cognition is an abstract representation of the way people make decisions and generate behavior. Much has been written and much has yet to be written about what models describe “true” cognition. For a person who designs or evaluates human-machine systems, the resulting debates provide fruitful ground for identifying those principles that are relevant for human machine interaction. In this paper, we present a model of cognitive information processing that we have used extensively in designing experiments in the field of human robot interaction (HRI). This model combines aspects of several models in the literature, and represents those components that we have found most useful. Although the model does not represent all aspects of cognition (no model does), it has proven useful in guiding system designs and performing system evaluations. After presenting the model, we first discuss some design principles based on the model, and then review some experiments that have used components of the model to predict and describe outcomes.

A Model of Cognition

The risk of creating a model of human information processing is that it is too abstract to describe all cognitive processes

*Partial funding for this research was provided by DARPA under grant number NBCH1020013, and by Nissan Motor Company. Copyright © 2004, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

and may therefore be limited in what phenomena it can predict or eliminate. We believe that the model we have constructed has the following characteristics. The model is

- Cognitively plausible. Humans generate skilled behavior “as if” behavior was generated by such a model.
- Computationally practicable. The model requires only computations that can be performed by neural machinery.
- Perceptually feasible. Behavior is generated only by percepts that can be obtained from the environment.

This model is not a perfect description of cognition, but it does allow us to focus experiments on combining cognition and skilled performance.

Perceptual Attention, Response Selection, and Response Execution

In our work on designing and evaluating HRI systems, it has been helpful to create a model of human information processing. We have constructed such a model from the current literature on attention and working memory. This model is diagrammed in Figure 1. Beginning at the left side of the figure, environmental stimuli may be allowed into sensory short-term memory (SSTM). SSTM is typified by iconic memory from human vision, but there are multiple such memories depending on the category of environmental stimuli. At the minimum, there appears to be a channel for visual stimuli, a channel for auditory stimuli, and a channel for haptic stimuli (Wickens & Hollands 2000). There are a limited number of stimuli that can receive *attention*. We adopt a *gating* interpretation of how attention can restrict what stimuli enter SSTM, and use gating to explain how focus of attention can limit which stimuli receive attention. To distinguish the type of *attention* associated with gating from the type of attention that limits response selection, we refer to the former as *perceptual attention* and the latter as *response-selection attention* (Pashler 1997).

Although there is a lot of information that can be stored in SSTM, not all such information is useful for skilled behavior. More specifically, we adopt the perspective of modern modal descriptions of memory which suggests that there exists a short-term memory (STM) used to further process environmental stimuli. To help understand the various roles of STM and SSTM, it is helpful to note that SSTM is a

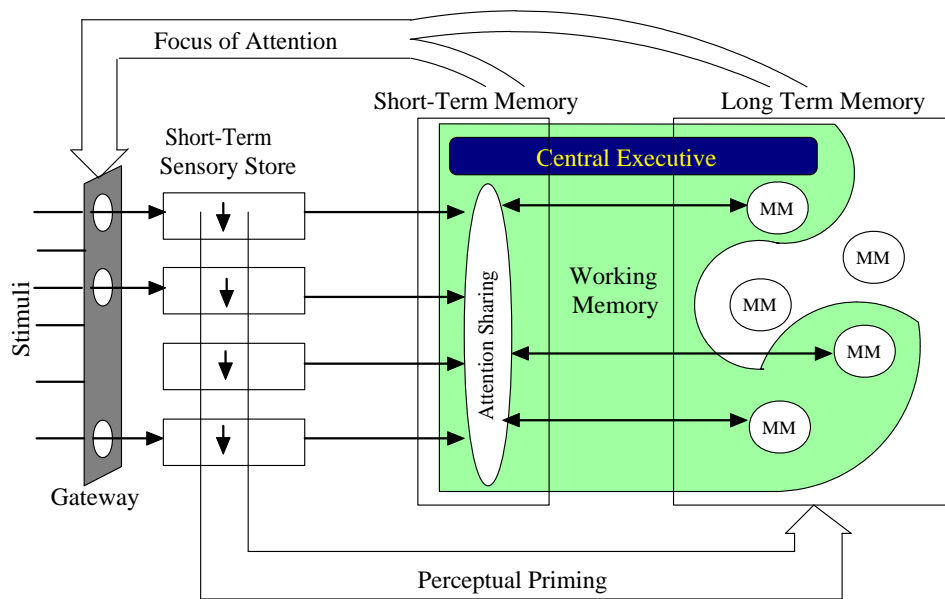


Figure 1: Interrelationships between computational elements.

wide-channel, extremely short-term memory and that, consequently, not all stimuli stored in SSTM are used to generate responses. Stimuli that are relevant for generating responses may be further processed by shifting them into a temporary memory store that acts as a scratch pad for further processing. For example, while driving an automobile, iconic memory may store something like a snapshot containing a huge amount of visual information, but only those visual cues (such as time-to-contact and time-to-lane-crossing) relevant for generating responses are transferred to short-term memory. Which items are transferred from SSTM to STM depends on which behaviors are currently receiving response-selection attention. This interpretation is consistent with a depth-of-processing interpretation that suggests that not all stimuli are processed to the same depth. Furthermore, this interpretation seems consistent with recent results that suggest that drivers who attend to, for example, cell-phones (or other distracting mental computations) increase the frequency of saccades but decrease the dwell time on each stimuli (Lee 2002); such drivers “see” more but perceive less.

Working memory is a theoretical construct generated by cognitive psychologists to encapsulate phenomena associated with restricted information processing. Although many interpretations of working memory exist (Miyake & Shah 1999), we note that many of these interpretations seem compatible with modal descriptions of short-term memory, and at least some of these interpretations insist on a component of long-term memory (Ericsson & Delaney 1999). We note that interpretations that rely on long-term memory are remarkably similar to descriptions of mental-models (Johnson-Laird 1988). Thus, we adopt the perspective that working memory consists of *information* stored in short-term memory and *processes* encoded as mental mod-

els in long-term memory. This has obvious associations with declarative and procedural knowledge, but we view a mental model as the unit that combines these two forms of knowledge into a behavioral quanta and short-term memory as a “scratch” pad for percepts needed by a mental model.

Since people successfully perform multiple tasks, we adopt the stance that multiple mental models can be concurrent in working memory. However, we also adopt Pashler’s bottleneck theory that suggests an information-processing pipeline where only one response may be generated at a time (Pashler 1997). Thus, there may be multiple mental models concurrent in working memory, but only one mental model can generate a response at a time. Furthermore, not all possible mental models can be concurrent at the same time; the number of active mental models is limited by the bounds of working memory.

To coordinate which mental model uses response-selection module at a given time, we adopt the stance that there exists a special mental model, called the *central executive*, that coordinates the other mental models (Baddeley 1986; Shallice 1988). This mental model shares response-selection attention with the other mental models and acts to help schedule response-selection attention. The central executive is a convenient fiction; the coordination of multiple mental models is probably better described by neural activation models, but using the neural level of detail does not contribute much to the kinds of problems that we are interested in solving.

The elements of working memory guide the focus of perceptual attention by inhibiting the processing of irrelevant stimuli and enhancing the processing of relevant stimuli. Thus, our model depicts a control mechanism from working memory to the perceptual gateway. This causes task-relevant stimuli to receive precedence over task-irrelevant stimuli.

Fortunately, there also exists a control mechanism that associates stimuli in SSTM with elements in working memory. Such a mechanism, known as *priming*, allows some mental models to become resident in working memory even before response-selection attention is allocated to them. In other words, a salient stimuli may cause a relevant mental model to enter working memory even if not explicitly instructed to do so by a context-sensitive central executive. This allows for certain stimuli to pop-out to a human even if unexpected, and may allow stimuli obtained from redundant channels to prevent an important mental model from being expelled from working memory.

Working Memory and Multiple Mental Models

Working memory, as we have modelled it, subsumes short term memory and includes certain active mental models from long-term memory. Before giving further clarification of how mental models can be active, it is useful to be more precise about what is meant by a mental model.




Figure 2: Working specification of a mental model.

A mental model is an internal representation employed to encode, predict, and evaluate the consequences of perceived and intended changes to the operator’s current state within a dynamic environment. A mental model is, in essence, a “chunk” of memory used to modify data in short term memory or to map data from short term memory to behavior.

Formally, we define a mental model \mathcal{M} as a triple consisting of the perceived state of the environment S , a set of decisions or actions A , and a set of ordered consequences C that result from choosing $a \in A$ when the environment is in state $s \in S$. According to this specification, a mental model not only encodes the relation between input s , action a , and perceived consequence c , but also includes both a notion of preferences among consequences as well as a notion of the frequency with which events (e.g., consequences) occur (Wickens & Hollands 2000, Page 72) (see Figure 2, and compare to related figures in (Meystel 1996; Sheridan 1992; Albus 1991)).

Each mental model can be categorized using Rasmussen’s knowledge-based (KB), rule-based (RB), or skill-based (SB) categories (Rasmussen 1976; Sheridan 1992). KB mental models must rely on general computing mechanisms to process stimuli and generate behaviors. They likely consume more response-selection attention and more STM than more skilled behaviors. RB mental models must rely on heuristics

and other explicit condition-action rules to generate behaviors. They place a greater demand on working memory by requiring rules to be stored in long-term memory and evaluated in short-term memory. The search through the space of possible rules consumes response-selection attention, but the execution of the behavior may be very fast. SB mental models require perceptual attention and likely require some response-selection attention¹, but the use of such resources is minimal. Many human behaviors become more *automatic* as they are practiced (Zsombok & Klein 1997; Simon 1996); increasing automaticity consists of movement from KB to RB to SB. In general, as behaviors become practiced, they become more skill-based and therefore require less work by the human.

It is now important to identify how working memory can include a subset of possible multiple mental models from long-term memory. Each mental model \mathcal{M} in long term memory will be described as being enabled/disabled and engaged/disengaged depending on whether it is influencing behavior generation and consuming response-selection attention, respectively. When \mathcal{M} is *enabled* the mental model is actively influencing human behavior generation, that is, it is taking a turn with the mechanism(s) responsible for causing response-selection bottleneck; and when *disabled* the mental model has no direct influence upon behavior. Although it is possible for humans to work “open-loop” by selecting behaviors with very little data, we assume that enabled mental models utilize perceptual resources a la perceptual attention.

When a mental model is *engaged*, the mental model holds relevant information in short-term memory whence environmental information is actively perceived and interpreted, and when *disengaged* the mental model releases its claim to short-term memory resources. This means that a mental model may be engaged (i.e., in working memory) even if it is not currently generating a response (enabled).

In terms of Figure 2, the mental model is enabled if the arcs between the mental model and behavior/actuation are active (whence behavior a is actuated) and the mental model is engaged if the arcs between the mental model and sensor/perception are active (whence s is actively perceived). We suppose that \mathcal{M} need not be enabled to be engaged, but an enabled mental model must at least time share response-selection attention. We have not modelled a structure that manages which mental models contribute to behavior generation and which consume attentional resources. Rather, we have tried to identify those design elements that help humans have the right mental model at the right time.

Design Principles and the Model

In a previous paper (Goodrich & Olsen 2003), we presented a partial list of principles that apply to designing human-robot systems. These principles were:

1. implicitly switch interfaces and autonomy modes,
2. let the robot use natural human cues,
3. manipulate the world instead of the robot,

¹This description of automaticity still requires attention and uses space in working memory. As such, it depends on Pashler’s behavior-generation description that minimizes short-term memory use and response-selection attention (Pashler 1997).

4. manipulate the relationship between the robot and world,
5. let people manipulate presented information,
6. externalize memory, and
7. help people manage attention.

To this list, we add

8. learn.

In the first draft of that paper, each of these principles was motivated by reference to the cognitive model, but these references were omitted in the interest of space in the final revision. In the remainder of this section, we briefly describe these principles and relate them to the cognitive model presented in the previous section.

Implicitly switch interfaces and autonomy modes.

Consider a frequently encountered HRI system that allows the operator to either enter waypoints on a map or teleoperate via a video feed. The obvious interface to this system would require the operator to explicitly select which mode they are using by, for example, a pull-down menu or a button. This requires that the user maintains more mental models than necessary; in addition to the *teleoperate* or *waypoint* mental model, the human must also keep a mental model engaged that tells them how to interact with the interface to switch between control modes. We refer to this latter mental model as the *interface management* mental model. It may be necessary for humans to have an interface management mental model, but we generally would like to minimize its role in working memory (to make the interface “transparent” — see, for example (Wren & Reynolds 2002)). In our work, we allow context to dictate which control mode is used; if the human grabs the joystick, the interface and autonomy mode automatically switch to support teleoperation, and if the human clicks on the map then the interface and autonomy mode automatically switch to support waypoint control. Since using the joystick or mouse in this way is already part of the human’s mental model, we eliminate the need for an interface-management mental model. The cognitive information processing model that we use predicts some resulting benefit in performance because superfluous mental models restrict how many task-relevant behaviors can be simultaneously managed.

Let the robot use natural human cues. An example of a system that allows a human to use natural human cues is Olsen’s work on safe/unsafe driving². This work allows a human to specify places where a robot can go by letting the human color regions in a digital image blue if the region corresponds to a safe place and red if the region corresponds to an unsafe place. The interface then automatically creates an image-based classifier that the robot uses to avoid unsafe places by classifying regions in the video feed. Since people use visual cues of what is safe and unsafe to help them navigate through the world, it is natural for them to perform this classification for a robot. Using natural human cues exploits pre-existing mental models and does not require the human to create a new mental model for how the robot perceives the environment. Thus, the cognitive information processing model predicts improved human-robot interaction for such systems.

²This is currently unpublished work done in the Interactive Computing Everywhere Laboratory at Brigham Young University.

Manipulate the world instead of the robot. An example of a system that allows a human to manipulate the world instead of the robot is one where a human touches the video at the point where they want to know what is going on. The robot then automatically drives to that location. This interface allows the human to request information about the world directly, without understanding, for example, how the robot translates inputs into wheel movements. Since the fundamental purpose of HRI is to allow a human to accomplish a task in the world (and not to allow the human to interact with a robot), this principle eliminates the need for the human to keep mental models of (a) how the robot will work as well as (b) the task they want to accomplish. Instead, the human will need only a mental model of the task they want to accomplish.

Manipulate the relationship between the robot and world. An example of a system that allows a human to manipulate the relationship between a robot and the world is a PDA-based interface for flying unmanned air vehicles (UAVs). Rather than controlling, for example, pitch to increase altitude, the operator instead clicks on a representation of the UAV on the display and drags it to a new height; see Figure 3. UAV autonomy then selects an ap-



Figure 3: This PDA display presents the relationship between the UAV and the world.

propriate control law that brings the UAV to the desired altitude (Quigley, Goodrich, & Beard 2004). This allows the operator to ignore the mapping between UAV control surfaces and to focus instead on placing the UAV in a world-based reference frame that will allow the operator to accomplish the assigned task. In this case, the world-based reference frame is the height of the UAV above the ground. The operator still needs a mental model for how the task

can be accomplished by changing the pose of the UAV in the world, but experience has shown that this mental model is easily learned by novices whereas the mapping between control surfaces to robot pose is only present in expert fliers.

Let people manipulate presented information. Many conventional robots include a camera, range sensors, a compass, and perhaps a GPS. Typical interfaces display data from these sensors side by side. In experiments in teleoperation, we have found that navigating through a cluttered course without hitting an obstacle is sometimes more easily done by attending to only range sensors and almost completely ignoring the camera. An example of an interface that allows people to manipulate presented information is one where the robot can be controlled by clicking directly in the display of range sensors. The information necessary for avoiding obstacles is in this display, and allowing the operator to guide the robot via this display eliminates the need for the operator to translate range readings into another coordinate frame to allow the robot to be controlled. This reduces the need for translation mental models, and frees cognitive resources for other tasks.

Externalize memory. Many conventional robots include cameras and range sensors. Typical interfaces display camera imagery and a visual display range readings side by side. Typically, camera imagery has a limited field of view, so a teleoperator must integrate the range readings and camera into a representation of the pose of the robot in the world. This requires that the operator frequently sample both displays and remember the relationships between the displays,




Figure 4: This interface integrates range and camera information in a single display.

and this imposes a burden on short term memory and requires a mental model to integrate the information. Experiments suggest that this load is sufficiently high that an operator who is controlling the robot cannot simultaneously look for targets in the environment while teleoperating (Casper & Murphy 2002). An interface that integrates range and camera readings into a single display, shown in Figure 4 reduces the burden placed on short term memory and simplifies the integration mental model. We are currently conducting experiments to validate this hypothesis; initial results are en-

couraging (Ricks, Nielsen, & Goodrich 2004).

Help people manage attention. Experiments in automation and teleoperation have shown that a salient mental model can quickly receive all attention under conditions of stress or distraction. This causes other mental models to be deactivated which leads to their extinction from working memory. When the robot or human is in a situation where this can occur, it is useful to help the human keep task-appropriate mental models in working memory by bringing their attention to them. As a very simple example, if a salient distracter exists in a situation where an operator is guiding a robot, then this distracter might start receiving all attention. Giving the robot the ability to detect when it is stuck and signal the operator helps the operator to turn attention from the distract back to the primary control task. Since many interesting HRI problems take place in complex worlds, it is likely that many circumstances will have salient distracters. Consequently, it is useful to help operators manage their attention properly. Preliminary results have indicated that as the autonomy of the robot increases to support longer periods of robot neglect, attention management is useful in helping people manage the robot in a timely manner.

Learn. Humans have a wealth of mental models for solving various problems and accomplishing various tasks. Machine learning can be used at the interface or robot level to adapt system activity to match existing mental models. An example of an interface that can do this is a force feedback steering wheel that adapts the force profile to maximize safety and minimize impact on human comfort. People have well-defined mental models about how steering wheel movements translate into vehicle behaviors, and learning how force feedback triggers correct behaviors should respect these pre-existing mental models.

Evaluation Examples and the Model

In this section, we will discuss how the cognitive model motivates the use of secondary task studies for evaluating HRI systems. We will then briefly review examples of such studies, and interpret their results using the model. These studies include a study of autonomy-assisted automobile driving, a study of an ecological display for teleoperation, and a study of attention management aids.

Secondary Task Studies

It is desirable to design HRI systems that will work in real situations and natural settings. These situations are characterized by the need for people to accomplish multiple tasks in the presence of distractions. The presence of multiple tasks and distractions place a burden on working memory by crowding short term memory with “off-task” information and engaging many secondary mental models in long term memory.

Although it is impossible to anticipate every circumstance in which an HRI system will be used, it is possible to explore how multiple tasks and distractions affect the usefulness of an HRI system design by doing secondary task studies. Such studies require humans to achieve a robot-centered task while also accomplishing secondary tasks. The rate,

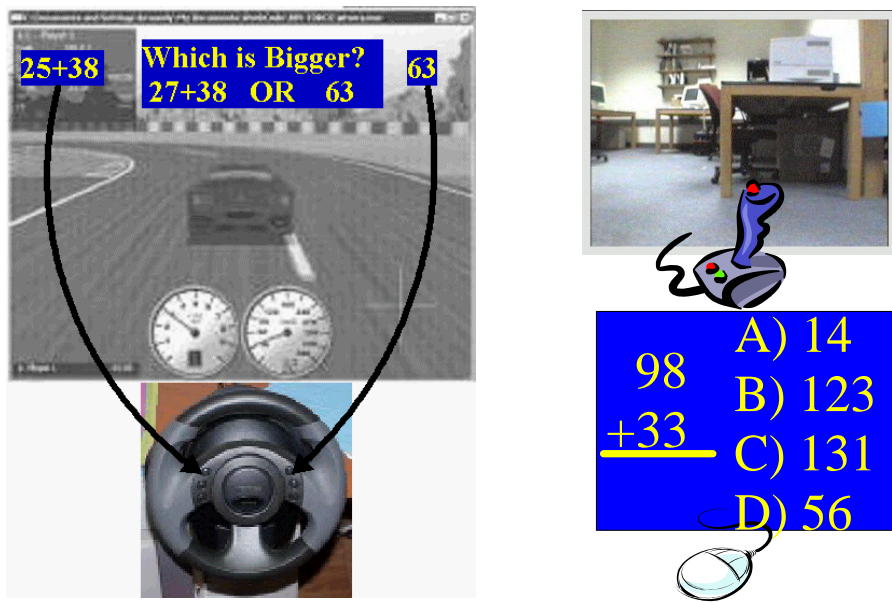


Figure 5: Examples of secondary task studies

difficulty, and sensor mode of these secondary tasks can be manipulated to explore various conditions of operator workload that might be encountered in the world.

Figure 5 illustrates two types of experiments that we have conducted: a driving task and an robot teleoperation task. In the driving task, subjects had to compare the result of an arithmetic problem to another number and then select the correct answer by pushing the appropriate button on the steering wheel; they had to do this task while staying in their lane and avoiding an erratic lead vehicle. In the robot task, subjects had to teleoperate via a video feed using a joystick while simultaneously selecting the correct answer to an arithmetic problem with their free hand.

In both types of experiments, the secondary task was a visual display of math problems that subjects were required to answer. The cognitive model suggests that if problems are presented visually they will divert visual attention from the primary task. Also according to the cognitive model, the math problems occupy space in working memory (partial answers in short term memory, and solving mental models in long term memory) and therefore interfere with the primary task. The answers were selected manually thereby affecting the manual control channels for steering wheel or the joystick. In the remainder of this section, we will summarize some observations from experiments that used this secondary task structure. This secondary task structure is a surrogate for doing studies in natural, distraction-rich settings.

Driving

A series of experiments were conducted to test the usefulness of force feedback in the steering wheel and gas pedal (Goodrich & Quigley 2004). The motivation behind these experiments was to see if communicating information

about lane position or lead traffic through the steering wheel and gas pedal, respectively, could improve driver response in natural driving scenarios. The cognitive model predicts that using the haptic channel to prime appropriate driving mental models even if the driver was distracted. Since the experiments were necessarily limited to a driving simulator, it was impossible to generate the true feel of natural driving including all possible distracters.

Instead, a secondary task study was performed. Subjects viewed math problems as they were displayed on the simulator screen at a controlled pace. (We also gave some math problems to them through headphones.) These secondary tasks simulated driver distracters. We then compared an optimized force feedback profile against nominal driving to see if the force feedback improved driver performance. The optimized force feedback profile was learned via a reinforcement learning algorithm (Goodrich & Quigley 2004) and gave corrective nudges through either the gas pedal or steering wheel.

For longitudinal (gas pedal) control, the benefit was clear. When a time headway value of 0.7 seconds was set as an “imminent danger” threshold, drivers spent 45% less time in the “imminent danger” zone with the haptic signal on the pedal versus without the signal. This large difference is the major advantage provided by the pedal forces.

Another metric which showed a consistent difference between trials was the average headway time between the subject and the lead car. With the pedal forces active, the average THW (across 7 formal test subjects) was 1.722 seconds, versus 1.676 seconds without the pedal forces. Despite the users’ overwhelming preference of the pedal forces, the average NASA TLX score only decreased from 70.65 to 70.47 indicating that the system increased safety without altering comfort; both average headway went up and minimum head-

way went up, but subjective workload estimates remained the same. These results were consistent with model predictions; the force feedback in the gas pedal helped people better schedule attention between the secondary task and the primary driving task.

By contrast, for lateral (steering wheel) control, the results depended heavily on the subjects. Subjects separated themselves into two categories: those that fought against the force feedback and those that admitted the forces. The subjective reports and objective data from the first category of subjects showed that their performance declined and their distaste for the force feedback was clear. By contrast, the subjective reports and objective data from the second category showed that their performance improved and showed that they tended to prefer the system.

We predicted that people would be better able to schedule attention, but the results showed otherwise. However, the results can be explained via the cognitive model. Simply put, some subjects had a strong mental model for how the steering wheel should feel, and this mental model was strong enough to prevent them from yielding to the corrective forces of the wheel. Other subjects were able to adopt a new mental model that allowed them to experience the benefits.

HRI Experiments

We have conducted a series of experiments on how robot autonomy, interface intelligence, and ecological displays affect people. Each of these experiments used the secondary task format to simulate natural conditions. The bottom line from each of these experiments is that if the people trust the system and understand how it works, then removing burdens on working memory through autonomy, intelligence, or display design improves people's ability to guide a robot.

Ecological Display Results. To summarize briefly, the ecological display integrates range and camera sensors into a single display; see Figure 4. Subjects guided a simulated robot through three mazes while performing a secondary memory recall task that loaded short term memory but did not interfere with visual attention or motor control. The cognitive model predicted that subjects would have fewer collisions and finish the mazes faster using this ecological display than with a display that presented range and video readings side-by-side. This prediction was supported by experiment results.

Interface and Robot Intelligence Results. To summarize the interface and robot intelligence experiments, human subjects were asked to solve secondary math tasks with and without robot path following and with and without an interface attention manager. These secondary math tasks were always present on the display and could be solved whenever subjects wanted to. The experiment predicted that path following would allow subjects to solve more secondary math tasks, but that attention management would be needed to help people re-attend to the robot. Experiment results support these predictions. When subjects had the path following accomplished the task faster and did more math problems than when they did not have it. However, there was a tendency for subjects to get "locked" into solving math prob-

lems and forget to re-attend to stuck robots. The attention manager significantly improved subjects' abilities to appropriately balance attention between the robot and the math problems.

Summary

We have reviewed a simple cognitive information processing model that has elements relevant to HRI. We then discussed how this model leads to a list of design principles, and how the model dictates the usefulness of secondary task studies for evaluating HRI system. The key elements of the cognitive model are the integration of short term and long term memory into working memory, and the role of mental models in generating task-appropriate behavior.

References

- Albus, J. S. 1991. Outline for a theory of intelligence. *IEEE Transactions on Systems, Man, and Cybernetics* 21(3):473–509.
- Baddeley, A. 1986. *Working Memory*, volume 11 of *Oxford Psychology Series*. Oxford.
- Casper, J., and Murphy, R. 2002. Workflow study on human-robot interaction in USAR. In *Proceedings of ICRA 2002, 1997–2003*. Finalist for the Anton Phillips Best Student Paper Award.
- Ericsson, K. A., and Delaney, P. F. 1999. Long-term working memory as an alternative to capacity models of working memory in everyday skilled performance. In Miyake, A., and Shah, P., eds., *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*. Cambridge University Press.
- Goodrich, M. A., and Olsen, D. R. 2003. Seven principles of efficient interaction. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, 3943–3948.
- Goodrich, M. A., and Quigley, M. 2004. Learning haptic feedback for guiding driver behavior. In *Proceedings of the 2004 International Conference on Systems, Man, and Cybernetics*. To appear.
- Johnson-Laird, P. N. 1988. *The Computer and the Mind: An Introduction to Cognitive Science*. Cambridge, Massachusetts: Harvard University Press.
- Lee, J. D. 2002. Eye movements under cognitive loading: Saccades and dwell-times. Personal communication of eye movement measurements made while drivers talked on a cell phone.
- Meystel, A. 1996. Intelligent systems: A semiotic perspective. In *Proceedings of the IEEE International Symposium on Intelligent Control*.
- Miyake, A., and Shah, P., eds. 1999. *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*. Cambridge University Press.
- Pashler, H. E. 1997. *The Psychology of Attention*. The MIT Press.
- Quigley, M.; Goodrich, M. A.; and Beard, R. W. 2004. Semi-autonomous human-UAV interfaces for fixed-wing mini-UAVs. In *Proceedings of IROS 2004*. To appear.
- Rasmussen, J. 1976. Outlines of a hybrid model of the process plant operator. In Sheridan, T. B., and Johansson, G., eds., *Monitoring Behavior and Supervisory Control*. Plenum. 371–383.
- Ricks, B.; Nielsen, C. W.; and Goodrich, M. A. 2004. Ecological displays for robot interaction: A new perspective. In *Proceedings of IROS 2004*. To appear.
- Shallice, T. 1988. *From Neuropsychology to Mental Structure*. Cambridge University Press.

- Sheridan, T. B. 1992. *Telerobotics, Automation, and Human Supervisory Control*. MIT Press.
- Simon, H. A. 1996. *The Sciences of the Artificial*. MIT Press, 3rd edition.
- Wickens, C. D., and Hollands, J. G. 2000. *Engineering Psychology and Human Performance*. Prentice Hall, third edition.
- Wren, C. R., and Reynolds, C. J. 2002. Parsimony and transparency in ubiquitous interface design. In *Proceedings of the 2002 International Conference on Ubiquitous Computing (UBICOMP)*, 31–33.
- Zsombok, C. E., and Klein, G., eds. 1997. *Naturalistic Decision Making*. Hillsdale, N.J.: Lawrence Erlbaum Associates.