Bennett, K. B. (1992). The use of on-line guidance, graphic displays, and discovery learning to improve the effectiveness of simulation training. In W. Regian & V. Shute (Eds.), Cognitive Approaches to Automated Instruction (pp. 217-244). Hillsdale, NJ: Lawrence Earlbaum Associates.

NOT A PAGE

9 The Use of On-Line Guidance, Representation Aiding, and Discovery Learning to Improve the Effectiveness of Simulation Training

Kevin Bennett
Wright State University, Dayton, Ohio

INTRODUCTION

Complex and dynamic real world domains (e.g., command and control, process control, air traffic control) require individuals to complete tasks that are difficult and involve significant risks. There is a great deal of interest in developing automated instructional systems that will allow an individual to acquire the required expertise without the associated risk and cost of training on actual equipment. One approach to automated instruction for these domains is the provision of a computerized simulation that replicates the domain of interest. A simulation provides one primary requirement for successful instruction: it allows an individual to acquire necessary skills and knowledge while actively solving domain problems. Simulations vary with respect to how well the physical appearance and dynamic behavior of a domain is replicated. In high-fidelity simulations the training system looks and acts just like the actual system. In low-fidelity simulations the emphasis is on providing instruction for the fundamental skills that underlay performance (Johnson, 1987). There are advantages and disadvantages to each end of the spectrum, and it is possible that instructional programs might be made more effective by incorporating simulations with varying levels of fidelity (Rouse, 1982).

One important issue in the design of simulation training environments is how to use the computational power that recent advances in computer science and artificial intelligence have provided. One approach is to develop and incorporate intelligent machine tutors. Intelligent tutors can provide advice, instruction, or explanation that dynamically changes as a function of the current situation and the individual's mastery of the topic at hand. With rare exceptions (e.g., Woolf, Blegen, & Jansen, & Verloop, 1987), intelligent tutors have not been incorporated into simulations for complex, dynamic domains. This is probably due to both the inherent complexity of these domains and the amount and types of knowledge (declarative knowledge, procedural knowledge, causal reasoning, psychomotor skills, and meta-cognitive skills) required for successful performance. Both of these factors create difficulties for the development of sophisticated student diagnosis and student modeling modules. One challenge that designers of simulation training environments face is how to incorporate this capability.

With the vast potential that intelligent machine tutors provide, it is easy to overlook other uses of computational power that can also improve the effectiveness of automated instruction. One technique is representation aiding (Hollan, Hutchins, & Weitzman, 1984; Woods & Roth, 1988), where machine power is used to create and manipulate representations of the domain. Direct representation and direct manipulation can help an individual find the relevant data in a dynamic environment, to visualize the semantics of the domain (i.e., make concrete the abstract), and to restructure his or her view of the problem. This can be an indispensable aid in facilitating understanding and performance. However, many issues must be resolved in the development of direct representation and direct manipulation interfaces. For example, which conceptual perspectives, of the nearly infinite number of alternatives, should be provided? What design perspectives are available to guide the mapping of information from the domain to the representation aids?

A third use of computational power is the provision of computerized tools and resources that allow discovery learning. In discovery learning, a student is encouraged to actively explore a domain, and computerized tools are provided to assist the student in the formation and testing of hypotheses about that domain (e.g., Shute & Glaser, 1990). The assumption is that actively exploring and experiencing the important domain concepts will result in more effective instruction than passively learning the same material. As with the development of intelligent tutors and representation aids, many questions arise with respect to discovery learning environments. Which of the potential explorations should be provided, and which computerized tools are needed to support these explorations? How should the explorations be sequenced to facilitate learning and transfer of training?

For the last few years, my colleagues (David Woods and Emily Roth) and I have been investigating issues in the design of decision support for complex and dynamic domains. We have chosen a representative domain (process control) and a representative task (the manual control of feedwater) for our investigations. We believe that on-line diagnosis, graphic displays, and discovery learning are three mutually reinforcing techniques that can be used to improve the effectiveness of training. We also believe that a detailed understanding of both the cognitive demands that are produced by the domain and the cognitive resources that skill domain practitioners have developed to meet these demands is required for

these techniques to be used appropriately. Two instructional systems have been developed that illustrate the approach. The first system is a part-task simulation that replicates the fundamental aspects that make the manual control of feedwater difficult. The second system is an on-line advisor that is incorporated within a full-scope simulator. These two systems, their development, and the design rationale are discussed.

THE MANUAL CONTROL OF FEEDWATER TASK AND THE PART-TASK SIMULATION

Any successful attempt to provide automated instruction needs to be based on cognitive analyses of the both the domain and expert performance within that domain. Providing this information is precisely the concern of a nascent discipline that has been referred to as "cognitive engineering" (Hollnagel & Woods, 1983; Norman, 1986; Rasmussen, 1986; Woods & Roth, 1988). Cognitive engineering provides a framework for the development of both on-line (real-time) and off-line (training) decision support. The approach can be paraphrased in the following manner. First, the "cognitive demands" that the underlying system places on the user must be determined. Then the user's "cognitive resources" (information processing capabilities, skills, knowledge, and higher-level strategies) that are available to meet these demands must be determined. Two mutually reinforcing, cognitive-based analyses were conducted to provide this information for the manual control of feedwater task (Roth & Woods, 1988; Woods & Roth, 1988). The goals/means analysis determined those aspects of a domain that make successful performance at the task difficult to achieve. The cognitive task analysis provided a description of the knowledge and skills that individuals had developed to overcome those difficulties.

Goals/Means Analysis

An analysis of the domain is crucial because the domain is the ultimate source of the cognitive demands that will be placed on an individual. As its name implies, the goals/means analysis provides a description of the domain in terms of goals that need to be accomplished and the physical resources that are available to accomplish them. Before discussing the details, a general description of the manual control of feedwater task is in order. The manual control of feedwater task (MCF task) is a critical and difficult task that can result in high economic losses when performance is poor. In a study of reactor trip data derived from plant monthly operation data over a 6-year period (1978–1983), it was found that the predominant cause of reactor trips was the manual control of feedwater task (INPO, 1984). This task was responsible for an average of 1.3 trips per plant, per year, and with certain plant configurations this estimate could be as high as five

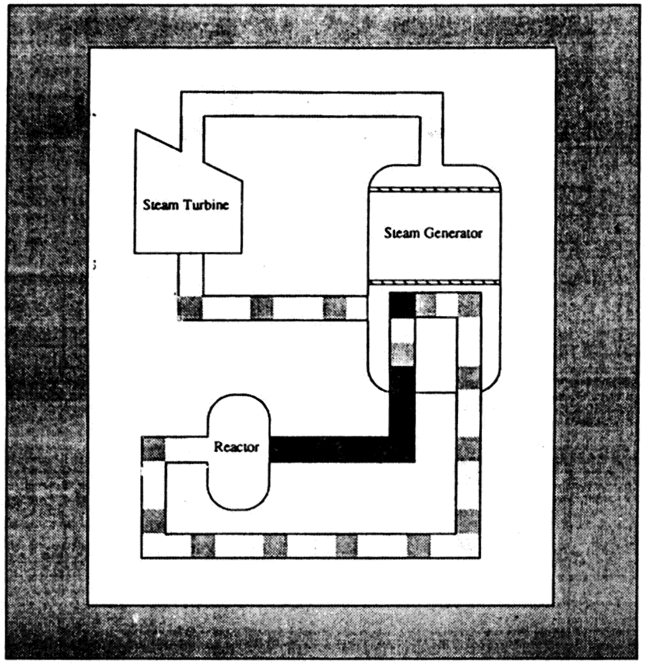

FIG. 9.1. The animated mimic display.

trips per plant per year. The task itself is embedded in the overall startup of a nuclear power plant and will be explained with reference to Fig. 9.1. During startup the energy produced by the nuclear reactor (in the form of heat) is piped through several boilers or steam generators (only one is pictured) to produce steam. This steam then drives a turbine to produce electricity and is subsequently cooled and returned to the steam generators in the form of feedwater. Controlling the rate at which feedwater returns to the steam generator is the manual control of feedwater task. In addition to the operator who performs this task, there are two more operators who control steam demand and reactor power. The critical performance variable is maintaining the level of water in the steam generators (the indicated steam generator level, ISGL) between an upper trip setpoint and a lower trip setpoint: crossing a setpoint automatically shuts the plant down and the startup must begin anew.

A goals/means analysis of the manual control of feedwater revealed that to accomplish the startup task two conflicting high-level goals must be achieved: (a) maintaining a mass balance by matching the amount of mass entering (feedwater flow) and leaving (steam flow) the steam generator, and (b) maintaining an energy balance by matching steam demand with reactor power production. As an example of how these goals can produce conflicts, consider the following scenario. Imagine that the ISGL is close to the lower setpoint boundary, and that, therefore, a high-priority goal is to raise this level to avoid a plant trip. One of the means available to accomplish this goal is to increase the rate of feedwater flow until the mass flowing into the steam generator is greater than the mass leaving it (steam flow). As one might expect, under these conditions the ISGL will ultimately rise. However, the increase of cold feedwater (relative to the environment inside the steam generator) initially decreases the thermal energy, causing the ISGL to fall (a "shrink" effect), and exacerbates the initial problem. This is one example of how energy effects can cause the initial change in ISGL to be the opposite of the long-term, steady-state effect and how ISGL reflects the true steam generator mass only after a time delay. The highly negative impact of even minimal time delays in tracking tasks is well documented (e.g., Wickens, 1986). Thus, the counterintuitive and time-delayed behavior of the ISGL is one fundamental source of difficulty in performing the task.

The goals/means analysis also revealed several other factors that contribute to the difficulty of the task. Each of the primary variables controlled by individual operators (steam flow, feedwater flow, reactor power) has an impact on the critical performance variable (indicated steam generator level). This high degree of intercoupling demands a significant amount of communication and cooperation between the three operators. In addition, the operators lack information about the critical primary variables: reliable estimates of steam flow and feedwater flow are not available at low power levels (where the startup task occurs) because plant sensors can only detect large volume flows. Finally, from two to four steam generators must be simultaneously controlled, adding a significant time-sharing component to the task.

Cognitive Task Analysis

Although the MCF task is quite difficult, expert operators have developed the knowledge and skills that allow them to accomplish the task. A cognitive task analysis was performed to determine the nature of those skills and knowledge. The cognitive task analysis defines the user's role within the goal/means hierarchy. That is, it determines the decisions and/or actions that the user must perform, and the data or information that is necessary to make those decisions or actions. Expert operators from nine separate utilities attended a 3-day panel session. The operators were asked to describe the major feedwater control maneuvers (with emphasis on what made these maneuvers particularly difficult) and incidents or near-incidents that they had been involved in or were familiar with, and they were also observed on a full-scope training simulator while actually performing the task.

The results of the cognitive task analysis indicated that expert operators had a detailed and rich understanding of the system that they were controlling. One facet of this understanding was the ability to recognize system state, involving the primary task of determining the extent to which current ISGL was due to long-term mass balances or transient energy effects. To accomplish this expert operators were observed to "mentally track" the influences that would eventually have an effect on ISGL, but that had not appeared yet due to the significant time delays. When expert operators lost track of these influences they were observed to perform "experiments"—entering small control inputs to observe the influence on ISGL. The cognitive task analysis revealed that a second facet of their detailed understanding of the system was the ability to predict ISGL response to control inputs. When mismatches between system goals and current system state (e.g., maintain the ISGL away from setpoints or maintain ISGL at a level that prepares for upcoming maneuvers) occur, then control input must be made. Expert operators were able to mentally simulate the system dynamics and choose between alternative control inputs that were most likely to achieve these goals.

In addition, expert operators exhibited specific control strategies that allowed them to avoid or to recover from problematic situations. One strategy to avoid trouble was to make small control inputs, thereby circumventing large oscillations in ISGL. Strategies to recover from trouble involved not only understanding the complex system dynamics, but actually using them to their advantage. Expert operators coordinated changes in the primary variables to produce artificial energy effects (referred to as shrink and swell levers) that would allow the required control inputs for recovery to be made. To recover from the scenario that was described earlier the feedwater operator might ask the turbine operator to increase steam demand, which would then cause a temporary increase in the ISGL (a swell effect), allowing the rate of feedwater to be increased without exceeding the lower setpoint boundary. Perhaps the most defining characteristic of an expert operator was the ability to quickly translate knowledge into action: expert operators were able to recognize a situation, select an appropriate strategy, and execute that strategy in an "automatic" fashion.

Part-Task Trainer

The findings of these two cognitive analyses guided the development of a part-task simulator for the manual control of feedwater task. A control theory expert familiar with nuclear power plants replicated the behavioral dynamics of a single steam generator with high functional fidelity. A set of differential equations were developed to reflect the influence of a number of factors on the ISGL, including steam flow, feedwater flow, and temperature of the feedwater. The simulation is generic in that the relative effects of these variables can be adjusted to represent a wide range of existing steam generators. Thus, the part-task trainer was designed to replicate the critical aspects of the task, as identified in the cognitive analyses described earlier. A significant amount of effort was devoted to developing additional decision support to replicate some of the critical skills and knowledge of expert operators.

Estimates of Steam and Feedwater Flows. The goals/means analysis and the cognitive task analysis revealed that one major contributor to the difficulty of the task is the lack of information regarding steam and feedwater flows. This information is critical because the relationship between these two variables determines mass balance. The mass flowing out of the steam generator (steam flow) must be replaced by mass flowing into the steam generator (feedwater flow) to maintain ISGL at a constant level. Although this information is not available in the actual plants, it can be obtained from the mathematical simulation and provided to individuals so that they can develop a deeper appreciation of the important relationship between them.

Compensated Steam Generator Level. The goals/means analysis revealed that two additional contributors to the difficulty of the MCF task are the long time delays between a control input and its effect on ISGL and the counterintuitive energy effects (shrink and swell). A prominent characteristic of expert performance was the ability to estimate the extent to which the current ISGL level was a result of these time delays and energy effects. A "compensated" steam generator level (CSGL) variable was developed that provides an estimate of the critical performance variable, indicated steam generator level. The CSGL variable eliminates the time lags and counterintuitive behaviors (shrink and swell effects) that are characteristic of the ISGL. When mass contributions are balanced (that is, when steam and feed flow are equal) the CSGL estimates the steady-state condition that ISGL will approach, and provides a direct indication of the size and direction of energy effects (shrink and swell).

Predicted Steam Generator Level. The cognitive task analysis also revealed that an important aspect of operators expertise was the ability to predict the behavior of ISGL. For instance, increasing the rate of feedwater flow could result in a shrink effect that would cause ISGL to cross the lower setpoint boundary. Normally the operators are required to mentally estimate the future behavior of the ISGL based on the current system context and their knowledge of the system dynamics. To assist the operators in this task, a predictor variable was developed that projects the value of ISGL into the future.

Summary

Developing effective automated instruction requires a detailed understanding of the target domain as well as the skills and knowledge that experts have developed to enable them to perform successfully within that domain. Cognitive engineering provides principles and techniques that can be used to discover this information. The analyses that were performed for the manual control of feedwater task provide one example of how the required information might be obtained (Mitchell & Saisi, 1987, describe a similar approach in the design of NASA ground control centers). The goals/means analysis determines the system goals that need to be accomplished, the physical resources (means) that are available to meet those goals, and the resulting cognitive demands that are placed on individuals performing tasks in the domain. The cognitive task analysis identifies the skills and knowledge (the cognitive resources) that individuals have developed to meet those demands. The ultimate goal is to identify instances where the cognitive demands of the domain of application and the cognitive resources of the user are mismatched. When the goal is to develop on-line decision support, these mismatches signal the need to reconfigure the existing interface (collecting or re-representing existing information) or the development of additional information to be added to the interface. When changes to the existing interface are not possible, cognitive demand/resource mismatches signal critical skills and knowledge that must be fostered by the instructional system. Having determined the critical information that is necessary for successful performance of the MCF task (variables, relationships between variables, goals and constraints), and having developed additional decision support to aid successful performance (estimates of steam and feed flows, the compensated SGL, the predicted SGL), the question becomes how to represent this information in the interface.

REPRESENTATION AIDING

One form of decision support that is often overlooked is representation aiding (Woods, 1991; Woods & Roth, 1988; Zachary, 1986) where machine power is used to create and manipulate graphic representations of the domain. Woods (Woods & Roth, 1988) and Rasmussen (1986) have stressed that there can be no neutral representation: any representation that is chosen will necessarily emphasize certain aspects of the domain at the expense of others. When designed appropriately, representation aids can be used to help the human problem solver find the relevant data in a dynamic environment, to visualize the semantics of the domain, and to restructure their view of the problem. This will be especially important during training and instruction, since an individual is explicitly learning about the domain. While technological developments have provided powerful capabilities to generate computer graphics, a clear understanding of how these capabilities can be used to support human cognition is needed.

There are several theoretical perspectives that can be used to guide the development of representation aids. Cleveland and his colleagues (Cleveland, 1985; Cleveland & McGill, 1985) have investigated the visual system's effectiveness in extracting information that has been mapped into various graphical forms (e.g., area of a circle vs. length of a line, etc.). Wickens and his colleagues (Wickens, 1986; Wickens & Andre, 1990; Wickens et al., 1985) have investigated the relationship between the general information-processing capabilities of an individual, the general demands of the task, and the implications for display design. Hutchins, Hollan, and Norman (1986) describe a general theory of interface

design that emphasizes the role of direct manipulation (the capability to effect changes in the domain by directly acting upon objects of interest).

A number of researchers have been investigating an alternative approach to display design for complex, dynamic domains. Although their theoretical orientations are slightly different, and the specific conclusions and recommendations may differ, they all share very similar basic beliefs. For these researchers, the success of a representation aid depends upon matching specific perceptual and cognitive capabilities of an individual with specific characteristics of the domain (Bennett, Toms, & Woods, in press); Flach & Vicente, 1989; Rasmussen, 1986; Vicente & Rasmussen, 1990; Woods & Roth, 1988). In particular, the semantics of those domains (the critical variables, the relationships between these variables, and the relevant goals and constraints) must be mapped into the static appearance and dynamic behavior of the representation aid so that critical information can be easily extracted or decoded by the individual.

PE

The principles that these researchers use to guide the design of representation aids for skilled performance are applicable to those for the acquisition of cognitive skills. However, the design of representation aids for automated instruction places additional requirements for integrated sets of displays. Representation aids also need to be designed to facilitate the transition from an initial understanding of the domain semantics to a more advanced conceptualization that approximates that of expert domain practitioners. In addition, when the graphic displays in the training system are not available on the target system, sets of representation aids need to be designed that facilitate the transfer of training to the target system. A set of representation aids that were developed for the part-task trainer is described in greater detail.

Trend Displays

In existing power plants, an operator performs the MCF task with a strip chart that provides the value of indicated steam generator level over time. Figure 9.2 illustrates two "trend" displays that were developed for the part-task stimulus. The upper trend display in Fig. 9.2 contains the critical performance variable ISGL and the CSGL. The current values of these variables are represented by diamonds in the right-hand portion of the screen; the history of these variables across a 5-min time frame is represented in the left hand portion of the screen. The lower trend display in Fig. 9.2 portrays the primary variables that effect ISGL level (steam flow, feed flow, and reactor power).

This representation of the domain semantics facilitates performance of the MCF task in several ways. The cognitive task analysis indicates that an important aspect of performing the MCF task is maintaining an internal record (over time) of the variables that influence ISGL. This is precisely the information conveyed by the trend displays. Variables that are normally available (the ISGL and power level) and some that are not normally available (steam flow, feed flow, CSGL, (the reactor core, the steam generator, and the turbine housing) and the pipes that connect them are represented graphically. The flow of steam (from steam generator to the turbine), feedwater (from turbine to steam generator), and energy (from reactor core to steam generator) is represented by animating the pipes. Apparent motion ~~animation~~ is produced by systematically changing (cycling) the luminance and chromaticity of adjacent squares inside the pipe. PE

Despite the intuitive appeal, very little empirical research has addressed issues associated with the implementation of this type of display. Basic research on the perception of motion indicates that the perceptual characteristics of the graphical elements (the squares) will have an impact on how well the apparent motion produced by the display matches the flow rates in the domain. These perceptual characteristics include fundamental frequency, temporal frequency, contrast (chromatic or luminance), shape, and borders. A series of empirical investigations have been conducted: the perceptual characteristics between the graphical elements were altered and observers matched the apparent motion of two horizontal, parallel pipes. The results of one set of experiments indicate that although chromatic contrast could be used to perform the rate-matching task, luminance contrast was much more effective. A second set of experiments revealed that there were optimal combinations of fundamental frequency (the size of the squares) and temporal frequency (the rate at which the squares were cycled). In a third set of experiments the nature of the borders between graphical elements was altered: the borders could be explicit (lines drawn between them) or implicit (no lines, just the contrast between elements) and the borders could remain vertical and fixed (no contours), or increasingly contoured (arrow-shaped) as rates of flow increased. It was found that explicit borders decreased rate-matching performance when no contours were present, but facilitated performance when the contours became more arrow-shaped. Additional details can be found in Bennett (1991a, 1991b).

The Configural Display

The animated mimic is an example of one type of display that has the potential to explain the physical functioning of complex causal systems. However, as Rasmussen (1986) has indicated, to accomplish tasks in complex domains the operator must understand the system at higher levels of abstraction, including system functions that cut across individual components or subsystems. Figure 9.3 illustrates a graphic display that has been developed to explicitly represent a higher-order, functional perspective of the MCF task. This representation corresponds to Rasmussen's fourth level in the hierarchy, the level of abstract function.

This type of display is referred to as a *configural* display. This display maps four variables (feedwater flow, steam flow, indicated steam generator level, and compensated steam generator level) into a single graphical object: a rectangle.

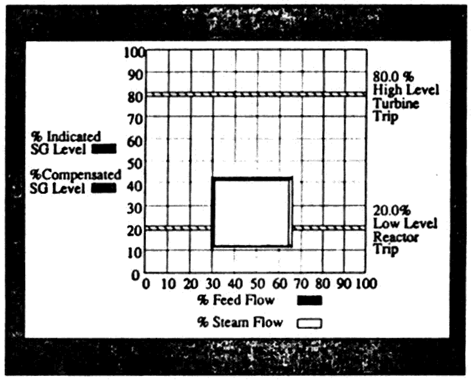

FIG. 9.3. The configural display.

The difference between steam and feedwater flow is mapped in the x axis (therefore, the width of the rectangle represents the degree to which mass input and output is balanced). The difference between the indicated and compensated steam generator level is mapped in the y axis (thus, the height of the rectangle roughly corresponds to the energy balance). The resulting display is a rectangle that changes in size and shape, as well as location inside the display grid. These changing perceptual cues provide direct information about the state of the plant. For example, a rectangle with a large area represents large imbalances in mass and energy and, therefore, an unstable and generally undesirable plant condition. This type of perceptual cue has been referred to as "emergent perceptual features" (Pomerantz, Sager, & Stoever, 1977; Sanderson, Flach, Buttigieg, & Casey, 1989) and ~~simply~~ does not exist with the display of the same information in a separable format (e.g., four bargraphs).

The success of a configural display in improving performance depends upon the extent to which the emergent perceptual features that result from the interaction of the lower-level graphic elements correspond to the demands of the task. When there is a direct correspondence, configural displays have been shown to facilitate performance (relative to separable displays) when information from several variables must be integrated to reach a decision (Wickens, 1986; Wickens & Andre, 1990; Wickens et al., 1985). The results of an empirical investigation that compared performance with the configural display in Fig. 9.3 to performance with a separable display of information (four bargraphs) support this conclusion (Bennett et al., in press). It was found that the configural display facilitated the accuracy (and in some conditions the latency) of responses that required the recall of critical task information: the mass and energy balances. Thus, the configural display provides a high-level conceptual perspective of the MCF task that emphasizes the functional, rather than the physical, aspects of the task that are critical for successful performance.

However, there is a potential cost associated with configural displays. Some research has indicated that displaying information in a configural format can decrease the availability of information ~~revealed~~ to individual variables. Thus, a critical issue in the design of configural displays is how to prevent this cost. Bennett et al. (in press) found that color-coding the lower-level, configural elements of the display partially offset these costs. Bennett and Flach (1991) review issues in the design of configural displays, including methodologies, associated patterns of experimental results, and relevant theories of design. They conclude that designing configural displays to allow the extraction of information related to both high-level properties and low-level data is a very distinct possibility.

related (PE)

Summary

Representation aiding has a vast potential to improve overall system performance in human-machine systems, due to both our impressive capabilities to process and utilize spatial information and the abundance of hardware and software to produce computer-generated graphics. Graphic displays can collect and integrate information, provide alternative conceptual perspectives, make the abstract concrete, and in some cases transform problem solving from a process that requires limited cognitive resources to one that capitalizes on virtually limitless perceptual resources. This section has outlined one design perspective that can be used to capitalize on this vast potential. One of the key conceptualizations is that the effectiveness of representation aiding depends on the mapping of information from the domain into the perceptual characteristics of the display. To the extent that the relevant information can be easily extracted from the representation, it will be effective.

The displays and decision support that have been developed for the part-task trainer were developed from this perspective. These displays should facilitate the acquisition of skills and knowledge required to complete the MCF task. The steam and feed flow variables and the compensated ISGL variable replicate a large portion of the expertise that expert domain practitioners exhibit. For example, these displays separate the effects of mass and energy balances on the critical performance variable ISGL and eliminate the long time delays associated with control input. Interacting with these displays should facilitate the development of skills such as determining system state, predicting future state, and anticipating necessary actions. Also, by eliminating the need to maintain an internal record of influences on ISGL, the demands on short-term memory are reduced, which allows an individual to concentrate on higher-level aspects of the control task such as control strategies to avoid or recover from trouble.

The trend displays, animated mimic, and configural display provide multiple conceptual perspectives of the domain semantics that pave the way for necessary transitions in knowledge. The animated mimic provides a very concrete, physical perspective of the system components and the causal relationships between them.

This display should be particularly appropriate for the development of an appropriate mental model of the MCF task and provide a basis for causal reasoning. On the other hand, the configural display provides a representation that is a higher-level abstraction of the functional relationships between the primary variables (mass and energy balances). This representation corresponds to one conceptualization that a more experienced operator may have developed after considerable experience at the task. Thus, these displays provide a framework for the transition from novice to expert conceptualizations of the task. In addition, the trend displays provide a conceptual perspective that matches the perspective in the target system and thus provides a means to transition the individuals from the additional decision support that is available on the part-task trainer that is currently available in the actual domain.

DISCOVERY LEARNING

One of the fundamental issues in automated instruction is the degree of learner versus system control of the instructional sequence (Glaser, 1990; Glaser & Bassok, 1989). One approach to automated instruction is to design systems that provide the learner with the opportunity to actively explore a domain. In this type of system a student is provided with tools that assist in the formation and testing of hypotheses about that domain (e.g., Shute & Glaser, 1990). This approach to instruction is based upon the belief that active participation in the learning process will improve the effectiveness of learning.

One of the primary resources that we have to support the discovery learning approach to instruction is representation aiding. As previous discussion has indicated, alternative conceptual perspectives and information that is not normally available can be provided to facilitate understanding and comprehension of complex domains. The discussion emphasized the importance of direct perception, that is, the representation of the domain semantics in a visual form that highlights the critical information necessary to accomplish domain tasks. However, as Holland and his colleagues (Hollan et al., 1984; Hutchins et al., 1986) have emphasized, direct manipulation can also be an important element of discovery learning approaches. Variables can be directly manipulated on the screen, and the resulting effects can provide immediate visual feedback to facilitate understanding.

Several lines of research support the potential of the discovery approach to facilitate learning. For example, Bobrow and Bower (1969) found that asking individuals to generate answers facilitated recall of sentences. Roth, Bennett, and Woods (1987) investigated the effectiveness of performance with an expert system that was designed to replace training on the troubleshooting of an electronic device. It was found that troubleshooting performance was best when technicians actively participated in troubleshooting activities, in direct contrast to the unsuccessful performance that resulted when technicians passively followed the instructions that were given by the system.

Lesh (1987) directly compared discovery and traditional learning approaches for the acquisition of mathematical skills. In subsequent tests, students experiencing the discovery learning approach performed better than those experiencing the traditional approach. In addition, individuals in the discovery learning condition were more likely to understand domain principles that were not directly addressed during instruction. It appears that actively exploring and experiencing concepts in the domain can result in increased structure and interconnections of knowledge in memory. The principles of discovery learning may provide one solution to the problem of inert knowledge (Bransford, Sherwood, Vye, & Rieser, 1986).

The combination of a discovery learning approach and graphic representations providing direct perception and manipulation have a great potential to improve automated instructional systems for complex, dynamic domains. Three key issues in the development of a discovery environment are (a) the nature of the explorations that are provided, (b) the nature of the tools and representations that are provided to support these explorations, and (c) the ordering or sequencing of these explorations. The goal of a discovery learning environment is to provide an individual with the capability to explore the complete range of the situations that might be encountered in a domain and the alternative behavioral responses that are available. A number of potential explorations that might be useful for the MCF task and other complex, dynamic domains will be discussed.

Change System Variables to Perform Experiments. This exploration is really at the heart of the discovery learning approach. In Smithtown (Shute & Glaser, 1990), individuals manipulate variables to perform experiments and their actions are monitored to see whether they adhere to the scientific method. One issue in this type of exploration concerns which variables should be manipulable. In the MCF task, individuals should be allowed to change any of the primary variables (steam flow, feed flow, and reactor power) or other variables (e.g., the temperature of the feedwater can vary, and this has a significant influence on the size of the shrink and swell effects that occur) that effect the ISGL. However, since ISGL is the critical performance variable that they must learn to control, and since ISGL is influenced by a number of variables, individuals should not be allowed to manipulate it.

Change System Displays. One important exploration is the capability to add or subtract additional decisional support, and to alternate between various conceptual perspectives. For example, interacting with the additional decision support and alternative conceptual perspectives is likely to facilitate the acquisition of skill at the MCF task. However, because this information will not be available during performance of the actual task, an individual needs the capability to add

or subtract this additional support. In this manner, individual can learn to perform the task with the additional decision support and then gradually "wean" themselves until they can perform the task with the information and conceptual perspective that is available in the real world.

Change Scenarios. In complex, dynamic domains one of the most important explorations that an individual should have is the ability to explore the range of situations that might be encountered. The scenarios should be carefully chosen to present situations that require all categories of the skills and knowledge that are necessary for successful performance in the domain, and ideally, instances within these categories that vary in their degree of difficulty. In the case of the MCF task, this might include standard scenarios with a low degree of difficulty, scenarios with existing trouble that requires the use of specific strategies to recover from trouble, scenarios with impending difficulty where strategies to avoid trouble are required, scenarios where other team members are working at cross purposes, and scenarios where critical aspects are varied (e.g., the temperature of the feedwater is varied to produce larger or smaller shrink/swell effects).

Repeat or Rewind a Particular Scenario. When learning to control a complex process, a learner will benefit from the opportunity to control the process again under the exact same circumstances. For instance, an individual who lost control and crossed a setpoint boundary in the MCF task should be able to immediately "rewind" the scenario and attempt an alternative strategy. Another version of this exploration is to attempt the same scenario with different decision support (e.g., successful completion of a scenario with the CSGL, followed by another attempt with ISGL alone).

Change Role. In complex, dynamic domains teams of individuals must often work together to achieve common goals. This requires an understanding of each team member's role in achieving overall goals. For example, the manual control of feedwater operator often coordinates the actions of the other two operators during the execution of strategies to avoid or recover from trouble (e.g., artificial shrink and swell effects). Thus, an important exploration might be the ability to perform complex tasks while assuming each team member's role.

Step Through or Speed Up Simulation. In complex, dynamic domains there are often multiple events occurring simultaneously, with the end result that an individual may not be able to attend to all events simultaneously. Allowing an individual to slow down the rate at which the simulation updates is an exploration that would this to happen. The converse of this exploration is time compression, which is the acceleration of the simulation update rate. In cases where there are significant time delays between a control input and the system's response, reducing the update rate can facilitate the acquisition of skill (e.g., Vidulich, Yeh, & Schneider, 1983).

Predict System Dynamics. In complex, dynamic domains, particularly those involving causal systems, one important exploration is to allow an individual to predict the behavior of the system or device (e.g., in the MCF task the critical prediction is the behavior of the ISGL). One exploration that might be provided is to allow an individual to make a prediction about the future behavior of the domain and then be provided with feedback regarding the appropriateness of the prediction.

Summary

This section has outlined the rationale behind a discovery learning approach to automated instruction, as well as a number of explorations that might incorporated into this type of system. BE (AA) The explorations should provide an individual with the opportunity to experience the range of situations that might be encountered in the domain as well as alternative response strategies that are available. Questions concerning how these explorations should be sequenced are closely related to fundamental issues in instructional theory. The strong view of discovery learning requires that an individual be allowed total freedom in the sequencing of explorations. However, there is some evidence that students low in ability or motivation may require explicit coaching to take advantage of this type of learning environment. In addition, there is no guarantee that students will explore all of the domain problem space. At the other end of the instruction spectrum is what has been referred to as the "mastery" approach (Glaser, 1990), where an individual has very little discretion over the sequencing of instructional treatments.

It is likely that an approach located between these two extremes will prove to be the most effective strategy for discovery learning. This has been referred to as a "guided discovery" approach. For example, Shute and Glaser (1990) allow students freedom in their choice of explorations until it is clear that they are not making progress toward instructional goals, at which point more directed instruction is provided. This approach could be extended by allowing an individual freedom within progressive levels of instruction (e.g., an individual might not be provided the opportunity to explore the utility of shrink and swell leverage strategies until success at standard control is demonstrated). Some of the sequencing decisions can be guided by the results of the goals/means analysis, the cognitive task analysis, or even intuition. However, a theory of learning that is adequately detailed to provide guidelines for the sequencing of explorations is needed.

COGNITIVE DIAGNOSIS FOR AUTOMATED INSTRUCTION

The discussion of how explorations should be sequenced serves as an introduction to the role of guidance in automated instructional systems. The approach to

instruction that has been discussed to this point emphasizes the role of direct perception, direct manipulation, and discovery learning. When coupled with a simulation of the domain this can provide an effective environment for the acquisition of skills and knowledge for complex, dynamic domains (Hollan, Hutchins, & Weitzman, 1984). However, as Anderson (1988) indicates, this is only part of the expertise that an automated instructional system should provide. An automated instructional system should also guide a student through the acquisition of domain skills, including the provision of advice or instruction in particularly difficult or novel situations, and the adaptation of the instructional sequence based on the student's current level of competency. What are the mechanisms behind the provision of this advice and what form (verbal, graphic) or "grain" (specific or general) should the advice take?

Sophisticated diagnostic and student modeling techniques have evolved to provide guidance in well-constrained domains (Anderson, 1988; VanLehn, 1988). Implementing these advanced techniques requires highly accurate psychological models (e.g., cognitive modeling; Anderson, 1988). With rare exceptions, these techniques have not been incorporated into instructional systems for complex, dynamic domains. In part, this is due to the complexity of the skills and knowledge that are required. Individuals must have detailed declarative knowledge about low-level, physical components of the domain and how they combine to provide higher-level functions of the domain. Individuals must be able to causally reason about the flow of information or resources between components and functions. They must also develop appropriate response strategies to avoid or recover from trouble, contextual knowledge about when these strategies apply, lower-level psychomotor skills to execute these strategies, and monitoring skills for the assessment of progress towards goals. Developing the highly accurate and detailed psychological models of performance that are required to implement more sophisticated diagnosis and modeling techniques would require a tremendous amount of effort (e.g., see Anderson, 1988, and VanLehn, 1988, for discussions concerning the difficulty of, and prospects for, developing diagnostic and modeling techniques for causal reasoning).

An additional aspect of complex domains that would appear to complicate the process of providing guidance is their dynamic nature. Scenarios develop over time, and new events can occur at indeterminate times, thus changing the nature of the problem to be solved. Therefore, the instructional system has the additional burden of assessing the evolving problem solving context to evaluate student actions. However, situation assessment can actually serve as a basis for the provision of instruction. Perhaps the best illustration is the work of Suchman (1987), who emphasizes the importance of "situated action" as the basis of instruction that is tailored for a particular instance and a particular individual. Combining situation assessment with simpler diagnostic techniques is one potential solution to the problem of providing automated instructional environments for complex, dynamic domains. An on-line advisory system was developed for the manual control of feedwater task (Roth, Woods, Elm, & Gallagher, 1987) that utilized situation assessment as one basis for the provision of on-line guidance.

On-Live Advisory System

The on-line advisory system was implemented on a full-scope simulation of a power plant and offers advice that is based on an analysis of expert performance at the MCF task. The design emphasis was to replicate the "cognitive competencies" that expert operators displayed. A portion of this cognitive competency (recognition of system state, the anticipation of future responses, and the separation of mass and energy effects on the ISGL) was provided by the decision support that was discussed earlier. In addition to this information, the on-line advisor provided advice with respect to the acceptable operating ranges for primary variables and also specific strategies for avoiding and recovering from trouble.

Advice on Acceptable Operating Ranges for Primary Variables. The findings of a cognitive task analysis of the MCF task indicated that expert operators maintained primary variables in a "comfort zone" that minimized the risk of crossing a setpoint boundary. The on-line advisor provides an estimate of the acceptable ranges for the primary variables (steam flow, feed flow, reactor power). One unique feature of this advice is that it is provided graphically, in the form of vertical bars (or operating bands) that are placed adjacent to the current values of the variables. These graphic operating bands change as a function of the current state of the system: when the plant is stable there is a wide range of acceptable values; when the plant is unstable a very narrow range of acceptable values would be provided.

Advice on Specific Strategies. The system also offers advice on specific strategies to avoid or recover from trouble. This advice takes the form of action scripts that describe the system state, the desired outcome (goal), and general recommendations for control input. The provision of this advice is triggered by an assessment of both current and future system state. The advice that is provided is based on taxonomies of specific plant situations and specific response strategies of operators (obtained from the cognitive task analysis). For example, consider an instance where the ISGL level was dangerously low due to a shrink effect and increasing the feedwater flow would result in crossing the lower setpoint boundary. The on-line advisor recognizes the situation from an assessment of both the current state of the system (ISGL and CSGL variables) and the future state of the system (the ISGL predictor). A response strategy that expert operators were observed to exhibit in this situation is then provided. Initially, an alphanumeric warning appears in an alarm window informing the operator of the adverse plant state ("shrink in progress"), and as the situation worsens the alarm

window would direct the operator to consider the advice presented in an action script window. In this example the advice is to create an artificial swell (the goal, "create swell," and the recommended control input, "increase steam flow") and then adjust the mass balance (the goal, "establish net inflow," and the recommended input, "increase feed flow").

Common Frame of Reference. One obvious feature of the on-line advisory system is that knowledge about the domain is incorporated into a "black-box" expert (Anderson, 1988; Burton & Brown, 1979): the expertise consists of computational algorithms with no attempt to replicate the psychological mechanisms involved. The use of a black-box expert avoids the cost and difficulty of detailed cognitive modeling, which, as Anderson (1988) has emphasized, is formidable even in well-constrained domains. One potential disadvantage of using a black-box expert in instructional systems is the problem of generating explanations acceptable to students. This problem is a particular instance of the more general problem of "opaque device" that has been discussed in the context of expert systems (Roth, Bennett, & Woods, 1987; Suchman, 1987) and computer systems in general (Brown, 1986).

The on-line advisor partially alleviates this problem by integrating advice and the rationale behind it to provide a "common frame of reference" or "mutual understanding" (Roth, Bennett, & Woods, 1987; Suchman, 1987). In the on-line advisor advice is provided as a function of both current situation assessments (ISGL and CSGL) and future situation assessment (the predicted ISGL). This advice is both graphic (recommended operating bands) and alphanumeric (action scripts) in nature, and the grain of this advice in general (high-level recommendations and not specific actions). In addition, the recommended control input is accompanied by a statement of the goal(s) for a particular maneuver. All of this information is directly available to the operator in the interface. This common frame of reference serves as both an explanation of the advice that is provided and a basis for evaluating the effectiveness of that advice.

Issue-Based Diagnosis

The major drawback in using situation assessment as the primary basis for instructional intervention is that instruction cannot be adapted according to an individual's current level of expertise or understanding. One technique that can be used to provide this capability is issue-based diagnosis (Burton & Brown, 1979). Issue-based diagnosis requires an analysis of domain tasks to determine "issues" or fundamental skills that are required for successful performance. The student's performance is monitored with respect to these issues during the process of solving domain problems. The diagnostic module looks for instances where the student has the opportunity to express the skills or knowledge associated with an issue, and maintains a record of whether or not this occurs. After sufficient behavioral evidence has accumulated with respect to an issue either advancement or instructional intervention can be provided. To implement the issue-based approach, the instructional system must be able to assess the evolving problem solving context, and have access to behavioral input within this context (what VanLehn, 1988, has referred to as "intermediate state" input). In addition, a model of expert performance is required for comparison purposes.

One positive feature of this approach to on-line cognitive diagnosis is the fact that detailed psychological models of all critical skills and knowledge may not be required for useful instruction to occur. Burton and Brown (1979, p. 8) state that "the black-box Expert used for evaluation need only be augmented with those incomplete pieces of an articulate Expert which are needed to detect critical or tutorable features of the answers produced by the black-box Expert. The glass-box Expert need not be able to produce the complete solution itself. It needs only to work backwards from the solution to determine the "important" (tutorial) features of the solution. This realization opens up the possibility of constructing coaching systems for domains for which we do not have complete glass-box expertise."

Although the development of the on-line advisor did not concentrate on cognitive diagnosis per se, many of the requirements for the implementation of issue-based diagnosis are present. The issues or fundamental skills that are necessary for the MCF task were identified in the goal/means and cognitive task analyses. Performance with respect to many of these issues can be assessed with direct behavioral measurements. In addition, the student's ability to select and execute appropriate control strategies could be measured by comparing the patterns and timing of their control inputs to those generated by the on-line advisor.

Summary

The MCF on-line advisor utilizes black-box expertise, provides advice that is based on analyses of expert performance, and uses situation assessment as the basis for the provision of this advice. Although this particular system has clear limitations, it provides one example of a broad theoretical approach to cognitive diagnosis (and instruction in general): guided discovery learning (Burton & Brown, 1979). From this theoretical perspective errors and mistakes are viewed as an important part of the learning process. For example, Burton and Brown (1979, p. 6) state that "While the student is making mistakes in the environment he is also experiencing the idea of learning from his mistakes and discovering the means to recover from his mistakes. If the Coach immediately points out the student's errors, there is a real danger that the student will never develop the necessary skills for examining his own behavior and looking for the causes of his own mistakes." There is a fair amount of data that supports the utility of this approach. Some of that data was reviewed in the section on guided discovery learning (e.g., Lesh, 1987; Shute & Glaser, 1990). The range of applicability for issue-based diagnosis is quite large: variations of issue-based diagnosis have been used successfully in a wide variety of domains and for a wide variety of knowledge (Anderson, 1988; Burton & Brown, 1979; Clancey, 1982). Issue-based diagnosis provides one way to cope with the complexity of the skills and knowledge that are required for successful performance in complex, dynamic domains. Although issue-based diagnosis relaxes the requirement to develop complete and highly-accurate psychological models, there is still a need to thoroughly understand the domain and the nature of the skills and knowledge that an individual uses to perform successfully in the domain. The cognitive task analysis and the goals/means analysis described earlier are examples of knowledge engineering methods that are required to support the approach.

SUMMARY

Cognitive diagnosis, student modeling, and other aspects of intelligent tutoring make up one way that instructional designers can use the abundant computational power that is currently available to improve the effectiveness of automated instruction. A second way to use this computational power is to develop representation aiding that allows individuals to envision the characteristics of complex domains through the provision of alternative conceptual perspectives (direct representation) and an improved capability to interact directly with the domain (direct manipulation). A third way to use this computational power is to provide aspects of discovery learning: explorations and computerized tools that support these explorations. Additional insight about how these three mutually interacting and supportive techniques can be used together will improve the effectiveness of automated instruction for complex, dynamic domains.

ACKNOWLEDGEMENTS

A number of individuals and organizations have contributed significantly to the ideas and systems that were discussed in the present paper. David Woods and John Flach provided helpful comments. David Woods also conceived and directed the development of the original part-task trainer and the on-line advisor (along with assistance from Emily Roth, Paul Haley, John Gallagher, Lew Hanes, and numerous other individuals at Westinghouse). Funding for the original development was provided by several Westinghouse organizations. Funding for additional research and development was provided by Wright State University (Research Challenge and Research Incentive Grants) and by the Air Force Office of Scientific Research (Summer Faculty Research Program and Research Initiation Program). A special thanks goes to Michael Young and his colleagues at AFHRL/LRG who hosted the SFRP appointment.

REFERENCES

Anderson, J. R. (1988). The expert module. In M. C. Polson & J. J. Richardson (Eds.), *Foundations of intelligent tutoring systems* (pp. 21–54). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bennett, K. B. (1991a). Representational aiding for a complex, dynamic control task. In *Proceedings of the 1991 IEEE International Conference on Systems, Man, and Cybernetics* (pp. 1207–1212). New York: IEEE.

Bennett, K. B. (1991b). *Encoding apparent motion in animated mimic displays*. Manuscript submitted for publication.

Bennett, K. B., & Flach, J. M. (1991). *Graphical displays: A search for meaning*. Manuscript submitted for publication.

Bennett, K. B., Toms, M. L., & Woods, D. D. (in press). Emergent features and configural elements: Designing more effective configural displays. *Human Factors*.

Bennett, K. B., Woods, D. D., Roth, E. M., & Haley, P. H. (1986). Predictor displays for complex, dynamic tasks: A preliminary investigation. In *Proceedings of the Human Factors Society 30th Annual Meeting* (pp. 684–688). Santa Monica, CA: Human Factors Society.

Brown, J. S. (1986). From cognitive to social ergonomics and beyond. In D. A. Norman & S. W. Draper (Eds.), *User centered system design* (pp. 457–486). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bobrow, S., & Bower, G. H. (1969). Comprehension and recall of sentences. *Journal of Experimental Psychology, 80,* 455–461.

Burton, R. R., & Brown, J. S. (1979). An investigation of computer coaching for informal learning activities. *International Journal of Man-Machine Studies, 11,* 5–24.

Clancey, W. J. (1982). Tutoring rules for guiding a case method dialogue. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems.* (pp. 201–225). New York: Academic Press.

Cleveland, W. S. (1985). *The elements of graphing data.* Belmont, CA: Wadsworth.

Cleveland, W. S., & McGill, R. (1985). Graphical perception and graphical methods for analyzing scientific data. *Science, 229,* 828–833.

Flach, J. M., & Vicente, K. J. (1989). *Complexity, difficulty, direct manipulation and direct perception* (Tech. Report EPRL-89-03). Urbana-Champaign: Engineering Psychology Research Laboratory, University of Illinois.

Glaser, R. (1990). The reemergence of learning theory within instructional research. *American Psychologist 45*(1), 29–39.

Glaser, R., & Bassok, M. (1989). Learning theory and the study of instruction. *Annual Review of Psychology, 40,* 631–666.

Hollan, J. D., Hutchins, E. L., & Weitzman, L. (1984). Steamer: An interactive inspectable simulation-based training system. *AI Magazine, Summer,* 15–27.

Hollnagel, E., & Woods, D. D. (1983). Cognitive systems engineering: New wine in new bottles. *International Journal of Man-Machine Studies, 18,* 583–600.

Hutchins, E. L., Hollan, J. D., & Norman, D. A. (1986). Direct manipulation interfaces. In D. A. Norman & S. W. Draper (Eds.), *User centered system design* (pp. 87–124). Hillsdale, NJ: Lawrence Erlbaum Associates.

INPO. (1984). Reactor trips caused by main feedwater control problems. *Significant Operating Experience Report, 84-4* (July).

Johnson, W. B. (1987). Development and evaluation of simulation-oriented computer-based instruction for diagnostic training. In W. B. Rouse (Ed.), *Advances in man-machine systems research* (Vol. 3, pp. 99–127). Greenwich, CT: JAI Press.

Lesh, R. (1987). The evolution of problem representations in the presence of powerful conceptual amplifiers. In C. Janvier (Ed.), *Problems of representation in the teaching and learning of mathematics* (pp. 197–206). Hillsdale, NJ: Lawrence Erlbaum Associates.

Mitchell, C., & Saisi, D. (1987). Use of model-based qualitative icons and adaptive windows in workstations for supervisory control systems. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-17,* 573–593.

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

Pomerantz, J. R., Sager, L. C., & Stoever, R. J. (1977). Perception of wholes and of their component parts: Some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 422–435.

Rasmussen, J. (1986). *Information processing and human-machine interaction: An approach to cognitive engineering.* New York: North Holland.

Roth, E. M., Bennett, K. B., & Woods, D. D. (1987). Human interaction with an "intelligent" machine. *International Journal of Man-Machine Studies, 27,* 479–526.

Roth, E. M., & Woods, D. D. (1988). Aiding human performance: I. Cognitive analysis. *Le Travail Humain, 51*(1), 39–64.

Roth, E. M., Woods, D. D., Elm, W. C., & Gallagher, J. M. (1987). *Intelligent manual feedwater control station* (Tech. Rep. 87-1-C60-HUSCI-R1). Westinghouse Research & Development Center.

Rouse, W. B. (1982). A mixed-fidelity approach to technical training. *Journal of Educational Technology Systems, 11*(2), 103–115.

Sanderson, P. M., Flach, J. M., Buttigieg, M. A., & Casey, E. J. (1989). Object displays do not always support better integrated task performance. *Human Factors, 31*(2), 183–198.

Shute, V., & Glaser, R. (1990). A large-scale evaluation of an intelligent discovery world: Smithtown. *Interactive Learning Environments, 1,* 51–77.

Suchman, L. A. (1987). *Plans and situated actions: The problem of human machine communication.* New York: Cambridge University Press.

VanLehn, K. (1988). Student modeling. In M. C. Polson & J. J. Richardson (Eds.), *Foundations of intelligent tutoring systems* (pp. 55–78). Hillsdale, NJ: Lawrence Erlbaum Associates.

Vicente, K. J., & Rasmussen, J. (1990). *The ecology of human-machine systems II: Mediating "direct perception" in complex work domains* (Tech. Report EPRL-90-01). Urbana-Champaign: Engineering Psychology Research Laboratory, University of Illinois.

Vidulich, M., Yeh, Y. Y., & Schneider, W. (1983). Time-compressed components for air-intercept control skills. *Proceedings of the 27th Annual Meeting of the Human Factors Society* (pp. 161–164).

Wickens, C. D. (1986). *The object display: Principles and a review of experimental findings* (Tech. Rep. CPL-86-6). Urbana-Champaign: Cognitive Psychophysiology Laboratory, University of Illinois.

Wickens, C. D., & Andre, A. D. (1990). Proximity compatibility and information display: Effects of color, space, and objectness on information integration. *Human Factors, 32,* 61–78.

Wickens, C. D., Kramer, A., Barnett, B., Carswell, M., Fracker, L., Goettl, B., & Harwood, K. (1985). *Display-cognitive interface: The effect of information integration requirements on display formatting for C3 displays* (Tech. Rep. EPL-85-3). Urbana-Champaign: Engineering Psychology Research Laboratory and Aviation Research Laboratory, University of Illinois.

Woods, D. D. (1991). Representation aiding: A ten year retrospective. In *Proceedings of the 1991 IEEE International Conference on Systems, Man, and Cybernetics.* Charlottesville, VA: IEEE.

Woods, D. D., & Roth, E. M. (1988). Cognitive systems engineering. In M. Helander (Ed.), *Handbook of human-computer interaction* (pp. 1–41). New York: North-Holland.

Woolf, B., Blegen, D., Jansen, J. H., & Verloop, A. (1987). Teaching a complex industrial process. In R. Lawler & M. Yazdini (Eds.), *Artificial intelligence and education, Vol. 1, Learning environments and tutoring systems* (pp. 413–427). Norwood, NJ: Ablex.

Zachary, W. (1986). A cognitively based functional taxonomy of decision support techniques. *Human-Computer Interaction, 2,* 25–63.