CHAPTER 2

Human-Automation Interaction

By Thomas B. Sheridan & Raja Parasuraman

Automation does not mean humans are replaced; quite the opposite. Increasingly, humans are asked to interact with automation in complex and typically large-scale systems, including aircraft and air traffic control, nuclear power, manufacturing plants, military systems, homes, and hospitals. This is not an easy or error-free task for either the system designer or the human operator/automation supervisor, especially as computer technology becomes ever more sophisticated. This review outlines recent research and challenges in the area, including taxonomies and qualitative models of human-automation interaction; descriptions of automation-related accidents and studies of adaptive automation; and social, political, and ethical issues.

The technological revolution ushered in by the computer has dramatically affected many aspects of human activity—at work and at home, during travel, and while engaged in leisure pursuits. Even more radical changes are anticipated in the next decade as computers decrease in size and cost and increase in power, speed, and "intelligence."

These factors are responsible for much of the drive toward increased automation in the workplace and elsewhere. The economic benefits that automation can provide (or is perceived to offer) has motivated considerable research and development on the technical capabilities of automation, which have been amply documented in such diverse domains as aviation; manufacturing; medicine; road, rail, and maritime transportation; robotics; home and entertainment devices; and numerous others. Humans work with or are consumers of all these technologies. Consequently, understanding how human characteristics and limitations influence the use (or misuse) of automation and using such knowledge to better the design of automated systems have been the focus of considerable research over the past two decades (Bainbridge, 1983; Billings, 1997; Jamieson & Vicente, 2005; Parasuraman & Mouloua, 1996; Parasuraman & Riley, 1997; Rasmussen, 1986; Sarter, Woods & Billings, 1997; Sheridan, 1992a, 2002; Wickens & Hollands, 2000; E. L. Wiener & Curry, 1980).

In this chapter we discuss research on humans and automation. We do not provide a comprehensive review of the field but describe recent and seminal work on the topic. We begin by defining automation and describe taxonomies and qualitative models of human-automation interaction, including the supervisory control model, function allocation, and the concept of human-centered automation. We then discuss automation-related accidents and incidents associated with inadequate feedback about system states, misunderstanding of automation, and overreliance. Subsequently, we describe recent research on the role of trust and "etiquette" in human-automation performance.

Because empirical human performance studies can be guided by and their results better interpreted with quantitative models, we briefly describe some of these models. Research on adaptive and adaptable automation is described next, with discussion of applications to driving and air traffic control. We then take a look at the future and discuss what the impact of new automation technologies might be. We close by discussing some social, political, and ethical issues that arise in considering the relationship between humans and automation.

WHAT IS AUTOMATION?

The *Oxford English Dictionary* defines automation as follows: "1. automatic control of the manufacture of a product through a number of successive stages; 2. the application of automatic control to any branch of industry or science; 3. by extension, the use of electronic or mechanical devices to replace human labor."

The first use of the term *automation* is traceable to a 1952 *Scientific American* article. Today use of the term has grown beyond product manufacturing and is applied to automatic control and instrumentation for chemical and power plants, aircraft and air traffic control, automobiles, ships, space vehicles and robots, heating and air conditioning in buildings, business systems, medical devices, home appliances, military systems, and stand-alone computers, to name only a few examples. Thus the second meaning is still widely accepted, as is the third meaning when human labor means mental as well as physical labor.

Mental labor is of primary importance for automation today, at least in the developed nations. Computers that interpret inputs, record data, make decisions, or generate displays are now regarded as automation, including the sensors that go with them, even though in the strict sense none of these functions may be automatically controlled.

In the fullest contemporary sense, the term *automation* refers to

a. the mechanization and integration of the sensing of environmental variables (by artificial sensors),
b. data processing and decision making (by computers);
c. mechanical action (by motors or devices that apply forces on the environment), and/or
d. "information action" by communication of processed information to people.

Automation can refer to open-loop operation on the environment or closed-loop control. It can apply to dynamic processes that change gradually over many months or those that occur in milliseconds. Thus contemporary definitions of automation refer to the gamut of processes, from the sensing of the environment to actions taken on that environment (Moray, Inagaki, & Itoh, 2000; Parasuraman, Sheridan, & Wickens, 2000).

HUMAN-AUTOMATION INTERACTION: TAXONOMIES AND QUALITATIVE MODELS

The Meaning of Human-Automation Interaction

What does it mean for human and automation to interact? Humans can be totally passive benefactors of automation: They purchase and use goods manufactured by automation. They consume electrical power, water, and heating and vehicular fuels that automation played a large role in providing. But being a passive user in this sense is not really interacting with the automation per se.

What we mean by human-automation interaction is the circumstances in which people (a) specify to the automation (necessarily a computer of some sort) the task goals and constraints (do X but avoid doing Y) and trade-offs between the goals and constraints; (b) control the automation to start or stop or modify the automatic task execution; and (c) receive from the automation information, energy, physical objects, or substances.

Simple examples of item 1 are people pushing the floor button on an elevator or setting the controls on their washing machines, and people setting the speed controls on their automobile cruise control system. More complex examples of item 1 include

a. pilots programming their flight management systems using a digital keypad or using special command language to have the autopilot take them to a new altitude and heading, fly to a series of designated waypoints, and enter the landing pattern at a distant airport;
b. machinists similarly programming a numerically controlled machine tool to make a metal part in a series of machining operations; and
c. space engineers programming movements of a robot arm on a Mars rover.

The conditions for the automation to start or stop or modify its program in the second item may be, for example, clock time or when sensors in a chemical plant indicate that a certain temperature has been reached, or when a robot has made contact with an object. Operators may abort automatic execution, and/or human takeover may be imposed when the human perceives a problem.

Much automation these days is used to give the operator information, either by warning or alarm display or by expert system to give advice. The information can be used by the human operator to reach a decision about some aspect of the system and to take action if necessary. But as in the case of decision aids, the automation may also provide a recommended choice or course of action. As stated in the third part of our definition, automation outputs can be energy (to move the whole body, as does an automobile or aircraft, or a part of the body, as does a prosthetic arm), an object (as with a vending or automatic teller machine), or a substance (as with an automatic drug delivery machine).

Supervisory Control Paradigm

The central role for humans in automated systems is to undertake what is called *supervisory control.* This is a new relation between the human and the machine, as an automatic machine may be said to be intelligent in some rudimentary sense. The new form

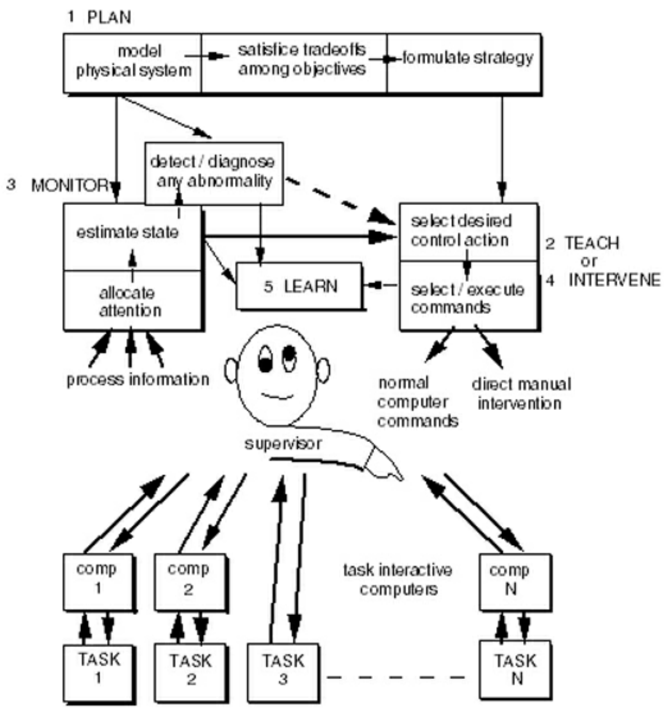

Figure 2.1. *Generic functions of supervisory control.*

of interaction differs dramatically from the traditional interaction of the human with tools and devices that possess no intelligence, in which all sensing and control were done by the human operator. This new relation was first called *human meta-control* (Sheridan, 1960) and later *human supervisory control* (Moray, 1986; Sheridan, 1992a Sheridan & Verplank, 1978). The human supervisor of the automatic machine was likened to the human supervisor of intelligent but subordinate humans, whereby the supervisor issued instructions (goals, constraints, and plans) and the subordinates executed those instructions using their own memories, built-in programs, sensors, and energy sources.

Figure 2.1 shows the relation of the human supervisor to the (typically multiple) computer-controlled machines performing simultaneous tasks (shown below the human figure). The several “task” boxes represent what might be different systems within a factory or power plant or aircraft, different degrees of freedom of a single robot, or multiple robots or unmanned vehicles the human supervises.

Above the human figure are boxes representing the five functions of the supervisor:

1. plan off-line,
2. teach the automation,
3. monitor the automation’s execution of the plan,
4. intervene to abort or assume control as necessary, and
5. learn from experience.

To plan involves having some mental or computer model of the physical system to be controlled, having some trade-off between performance objectives that can be satisficed (made acceptable), and formulation of a strategy for doing the task. To teach means to decide on a desired control action and to communicate commands to the automation to implement that action. To monitor means to allocate attention among the appropriate displays or other sources of information about task progress and from these to estimate the current state (vector) of the system (to maintain situation awareness). To intervene means that if an abnormality of sufficient magnitude is detected and diagnosed, the human will either reprogram the automation or may even take over and exercise manual control. Learning from knowledge of the results, like planning, is an out-of-the-loop human function and feeds back into planning the next phase of supervision.

In supervising a single or several parallel automated systems, as shown at the bottom of Figure 2.1, the human brings to bear whichever of these five functions is appropriate. At any instant of time the most appropriate function will typically differ from one task to another.

Sequential Stages of an Automated Large-Scale System

A large-scale system typically involves four classes of task, each a subtask or stage of a larger process (Figure 2.2):

a. the acquisition of information,
b. the analysis of that information,
c. the decision about actions to take based on that information, and
d. the implementation of that action.

Automation at each stage involves all five of the supervisor functions described earlier (a 4×5 matrix). For example, air traffic control involves (a) the acquisition of radar information on location, flight plans and identity of many aircraft, weather information, and so on. It requires (b) that appropriate information then be combined, analyzed, and displayed to the air traffic controller. Then it requires (c) that decisions be made as to speed, heading, and altitude for different aircraft to maintain safe separation and bring the aircraft safely though a sector of airspace or to land or take off. Finally, it requires (d) a means to get the pilots (and aircraft) to cooperate and execute the instructions given.

One can think of these four stages as the computer-controlled task boxes at the lower part of Figure 2.1. But they are not four independent tasks; they are four tasks, and each successive task depends totally on the previous one, as implied by Figure 2.2. These four stages can be automated, as Figure 2.1 implies. But they need not be fully automated. Some can be automated to a greater degree and others to a lesser degree (Parasuraman et al., 2000).

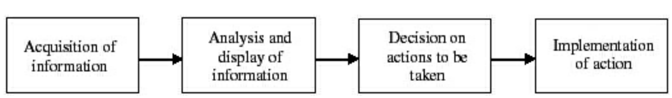

Figure 2.2. *Stages of a typical large-scale system.*

TABLE 2.1: A Scale of Degrees of Automation

1. The computer offers no assistance; the human must do it all.
2. The computer suggests alternative ways to do the task.
3. The computer selects one way to do the task and
4. executes that suggestion if the human approves, or
5. allows the human a restricted time to veto before automatic execution, or
6. executes the suggestion automatically, then necessarily informs the human, or
7. executes the suggestion automatically, then informs the human only if asked.
8. The computer selects the method, executes the task, and ignores the human.

Levels of Automation

Table 2.1 is a scale of degrees of automation. Alternative forms of such a scale are discussed in the literature (Endsley & Kaber, 1999; Sheridan, 1992b; Sheridan & Verplank, 1978; Wei, Macwan, & Wierenga, 1998). Some versions scale the sensing (afferent) and the motor (efferent) functions separately; some have more levels and some fewer. There are two main points to be made: that automation need not be all or none—there are various degrees appropriate to different problem contexts—and that different process stages of a complex system are appropriately automated to different degrees.

Parasuraman et al. (2000) emphasized the latter point and gave examples of how the appropriate level of automation differs at the four stages for different applications. Endsley and Kaber (1999) demonstrated the level of automation effects on performance, situation awareness, and workload in a dynamic control task. Wei et al. (1998) suggested a model for the appropriate degree of automation of different tasks based on a task's effect on system performance and its demand on the operator relative to other tasks.

Criteria of Function Allocation and Human-Centered Automation

Which functions should be relegated to the human and which to the automation is a classical problem going back to the Fitts (1951/2005) MABA-MABA list ("Men are better at . . . ; machines are better at . . ."). Although automation has progressed to the point that the Fitts list is no longer valid, criteria for function allocation have not been settled (Hancock & Scallen, 1996; Sheridan, 2000).

Human-centered automation is a phrase popularized by Billings (1997) and widely used to convey that automation must be designed to work in conjunction with the humans controlling or otherwise interacting with it; to engineer the automation and expect the human to accommodate to it can be a recipe for disaster. By now the point is well understood within the human factors community. But debate continues about

TABLE 2.2: Some Criteria of Human-Centered Automation (and Reasons to Question Them)

1. Allocate to the human the tasks best suited to the human, and allocate to the automation the tasks best suited to it. (Unfortunately, there is no consensus on how to do this; nor is the allocation policy necessarily fixed, but may depend on context.)
2. Keep the human operator in the decision-and-control loop. (This is good only for intermediate-bandwidth tasks. The human is too slow for high bandwidth and may fall asleep if bandwidth is too low.)
3. Maintain the human operator as the final authority over the automation. (Humans are poor monitors, and in some decisions it is better not to trust them; they are also poor decision makers when under time pressure and in complex situations.)
4. Make the human operator's job easier, more enjoyable, or more satisfying through friendly automation. (Operator ease, enjoyment, and satisfaction may be less important than system performance.)
5. Empower or enhance the human operator to the greatest extent possible through automation. (Power corrupts.)
6. Support trust by the human operator. (The human may come to overtrust the system.)
7. Give the operator computer-based advice about everything he or she should want to know. (The amount and complexity of information is likely to overwhelm the operator at exactly the worst time.)
8. Engineer the automation to reduce human error and minimize response variability. (A built-in margin for human error and experimentation helps the human learn and not become a robot; see Rasmussen, Pedersen, & Goodstein, 1995.)
9. Make the operator a supervisor of subordinate automatic control systems. (Sometimes straight manual control is better than supervisory control.)
10. Achieve the best combination of human and automatic control, where best is defined by explicit system objectives. (Rarely does a mathematical objective function exist.)

the criteria for whether automation is appropriately "human centered." As suggested by Table 2.2, the various criteria are different from one another, and for any one of these criteria, serious questions can be raised about the extent of its applicability.

AUTOMATION-RELATED INCIDENTS AND ACCIDENTS

Much of the impetus for understanding human interaction with automated systems stems from several accidents that have involved automation, an issue first raised in the aviation domain in a seminal paper by E. L. Wiener and Curry (1980). We focus on incidents and accidents related to feedback about system states provided by the automation, misunderstanding of automation, and overreliance on automation (Billings, 1997; Parasuraman & Byrne, 2003; Parasuraman & Riley, 1997). These three issues were also among the top 5 of more than 100 automation-related issues identified in a survey of aviation experts by Funk et al. (1999), the other 2 issues being poor display design and inadequate automation training.

Accident investigators have attempted to draw lessons learned from automation-related accidents, many of which have been in aviation (Billings, 1997), although accidents have also occurred in road, train, and maritime transportation; in manufacturing and process control; and in health care. Commercial aviation has a very good safety record, and modern highly automated aircraft are not only more fuel efficient but also safer than earlier generations of aircraft. Nevertheless, several highly publicized incidents and accidents involving automated aircraft in the 1980s and 1990s, coupled with the quest for even higher safety levels in the face of an increased volume of air traffic, have motivated greater scrutiny of automation by the aviation industry and the FAA (Abbott et al., 1996).

Although potential problems have been identified in human interaction with use of automation, it is not always easy to provide a succinct definition of an automation-related incident or accident because of the multiplicity of precipitating events and conditions ultimately leading to any accident (e.g., Reason, 1990; see also Wiegmann & Shappell, 1997). Accident investigators have used the National Transportation Safety Board Aviation Coding Manual (*Aviation Coding Manual,* 1998) and NASA's Aviation Safety Reporting System (Aviation Safety Reporting System, 2005), which can be examined with appropriate keywords to identify automated-related incidents (Funk et al., 1999).

Feedback on System States

Many so-called automation-related incidents occur even though the automation did not malfunction. Rather, in many incidents, the state of the automated system changed, but this was not communicated to the human operators in a salient way. An early example from aviation is the 1972 crash of a Lockheed L-1011 in the Florida Everglades. The flight crew was engaged in troubleshooting a problem with a landing gear indicator light and did not recognize that the altitude hold function of the autopilot had been inadvertently disconnected (National Transportation Safety Board, 1973). Although several factors contributed to this accident, a major one was poor feedback on the state of automation provided by the system. The disengagement of automation should be clearly signaled to the human operator so that it can be validated as intended or unintended. Most current autopilots now provide an aural and/or visual alert upon disconnect. The alert remains active for a few seconds or requires a second disconnect command input by the pilot before it is silenced. Persistent warnings such as these, especially when they require additional input from the pilot, are intended to decrease the chance of an autopilot disconnect or failure's going unnoticed.

However, a persistent warning may be insufficient by itself if the crew does not know what state the automation is in. As pointed out by Degani (2003) in his excellent recent book *Taming Hal,* many automated systems—from simple alarm clocks to automatic teller machines to large aircraft—involve internal transitions between different machine states or modes. Using state transition diagrams, Degani illustrated how such transitions are sometimes hidden from the user, as a result of which the user may think the machine is in one state when it is actually in another. With simple systems, such as VCR/TV controls, this might lead only to annoyance or frustration as the user fumbles with adjusting the TV set while the control is actually in VCR mode. But with more complex systems, the lack of salient feedback about automation states (Norman, 1990) can lead to catastrophe.

Two decades after the L-1011 accident, an Airbus A300 experienced an in-flight incident off the coast of Florida (National Transportation Safety Board, 1998b). At the start of a descent into the terminal area, the autothrottles were holding speed constant but no longer controlled airspeed after the plane leveled off at an intermediate altitude—a circumstance unknown to the pilots. The aircraft slowed gradually to almost 40 knots below the last airspeed set by the pilots and subsequently experienced stalling after the stall warning activated. There was no evidence of autothrottle malfunction. Rather, the crew apparently believed that the automated system was controlling airspeed when in fact it had disengaged, which could be done with a single press of the disconnect button. When the system was disengaged, the green mode annunciator in the primary flight display would change to amber, and the illuminated button on the glare shield used to engage the system would be extinguished.

The NTSB (1998b) noted that although the change in the annunciators could serve as a warning, the format of the displays did not command attention because they were passive and persistent. The NTSB also pointed to autothrottle disconnect warning systems in other aircraft that required positive crew action to silence or turn off. These systems incorporated flashing displays and, in some cases, aural alerts that would help capture the pilot's attention in the case of an inadvertent disconnect. These systems more rigorously adhere to the principle of providing salient feedback to the operator about the state of an automated system.

Misunderstanding or Lack of Understanding of Automation

A major characteristic of automated systems is complexity. Designers and engineers have developed automated systems with such large numbers of interacting subcomponents that understanding the effects of all possible interactions has become increasingly difficult, if not impossible. Much previous research has focused on misunderstandings of complex automated systems in the cockpit, such as the flight management system (FMS), whose complexity is so great that even highly skilled and trained operators such as commercial pilots can have difficulty understanding the nuances of its behavior (Sarter & Woods, 1995).

It has been suggested that misunderstandings arise because of a mismatch between the mental model of the pilot and the behaviors of the automated system as programmed by the designers (Sherry & Polson, 1999). Several examples of incidents and accidents resulting from these system misunderstandings have been reported (Billings, 1997; Funk et al., 1999; Sarter & Woods, 1995). Although some have had benign outcomes and become lessons learned, others have involved serious loss of life.

For example, in 1994, an A300 crashed in Nagoya, Japan, after the pilots inadvertently engaged the autopilot's go-around mode. The pilots countered the unexpected pitch-up by making manual inputs, which turned out to be ineffective (Billings, 1997). Essentially, the pilot attempted to continue the approach by manually deflecting the control column, which in all other aircraft—and in this aircraft in all modes except the approach mode—would normally disconnect the autopilot. However, in this particular aircraft and in this particular mode, the autopilot had to be manually deselected and could not be overridden by control column inputs. Consequently, a power struggle developed between the pilot and the autopilot, with the pilot attempting to push the nose down through elevator control and the autopilot attempting to lift the nose up through trim control. This caused the aircraft to become so far out of trim that it could no longer be controlled.

Overreliance on Automation

Automated systems typically are highly reliable—with the exception of some automated alerting systems, which can have high false alarm rates. This, together with their opacity and complexity, can lead operators to rely unquestioningly on automation. The phenomenon of overreliance on automation has been likened by Mosier, Skitka, Heers, and Burdick (1998) to a decision bias, or automation bias. They suggested that the bias is reflected in the operator's following the advice of automated systems even when the automation commits both errors of omission (misses) and errors of commission (false alarms).

However, *reliance* on automation can be distinguished from *compliance.* Meyer (2001) showed that when automation reliability is such that malfunctions are almost always correctly indicated, the automation makes few misses, so operators have high reliance on the automation. This is an effective strategy but can result in a problem when the automation *does* fail to indicate a hazard (a miss), because in this case the operator may not monitor the automation—the so-called complacency effect (Parasuraman et al., 1993). On the other hand, if automation reliability is such that few false alarms are made, then the operator usually has high compliance: If an automated alarm sounds, the operator tends to comply immediately with the alarm and attend to the situation. Reliance on automated aids permits the operator to attend to tasks other than the automated task until the alert is triggered, thus improving multitask performance and not just the performance on the automated task.

An example from the maritime industry is particularly revealing of the effects of user overreliance on automated systems. The cruise ship *Royal Majesty* ran aground off Nantucket after veering several miles off course into shallow waters. Fortunately, there were no injuries or fatalities as a result of the accident, but losses totaled $2 million in structural damage and $5 million in lost revenue. The automated systems in this ship included an autopilot and an Automatic Radar Plotting Aid that was tied to signals received by a Global Positioning System (GPS). Under normal operating conditions, the autopilot used GPS signals to keep the ship on its intended course. However, the GPS signals were lost when the cable from the antenna frayed (it was placed in an area of the ship where many sailors walked). As a result, the GPS and autopilot switched to dead reckoning mode, no longer correcting for winds and tides, which in this case carried the ship toward the shore.

According to the NTSB report on the accident, the probable cause was the crew's overreliance on the Automatic Radar Plotting Aid and managers' failure to ensure that crewmembers were adequately trained in understanding the automation features and its capabilities and limitations. The report went on to state, "the watch officers' monitoring of the status of the vessel's GPS was deficient throughout the voyage," and "all the watch-standing officers were overly reliant on the automated position display . . . and were, for all intents and purposes, sailing the map display instead of using navigation aids or lookout information."

This accident represents a classic case of automation complacency related to inappropriately high trust in the automation (Lee & See, 2004; see Degani, 2003, for a more detailed account of the accident). This accident also demonstrates the importance of salient feedback about automation states and actions, as mentioned earlier. The text annunciators that distinguished between the dead reckoning and satellite modes were not salient enough to draw the crew's attention to the problem.

A general aviation accident further exemplifies the danger of overreliance on automated systems such as GPS. In 1997, a single-engine airplane with a non-instrument-rated pilot took off under instrument meteorological conditions (National Transportation Safety Board, 1998a). About two hours later, following a meandering course that included course reversals and turns of more than 360°, the aircraft crashed into trees at the top of a ridge. No mechanical problems with the airplane's controls, engine, or flight instruments were identified. A person who spoke with the pilot before departure stated that the pilot "was anxious to get going. He felt he could get above the clouds. His GPS was working and he said as long as he kept the [attitude indicator] steady he'd be all right. He really felt he was going to get above the clouds." Many factors undoubtedly played a role in this accident, but the apparent reliance on GPS technology, perhaps to compensate for insufficient training and lack of ratings, stands out as a compelling factor.

There is some consensus for the existence of overreliance on automation (called *complacency*) for now-entrenched historical reasons related to the development of NASA's Aviation Safety Reporting System (see Billings, Lauber, Funkhouser, Lyman, & Huff, 1976). However, Moray and Inagaki (2001) disagreed with the interpretation of Parasuraman et al. (1993) that their findings reflected complacency. Though accepting that people have often failed to notice when automation fails, they argued that none of the reported literature has shown what the optimal level of trust would actually be. Consistent with the information-theoretic model of sampling (Senders, 1964), Moray and Inagaki (2001) proposed that automation should be monitored (or sampled) by the human at a rate set by the objective failure rate of the automation: The more reliable the automation, the less the operator should monitor it. Moray and Inagaki stated that complacency (overtrust) should be inferred only if the operator sampled less frequently than this rate; if they sampled more frequently, they should be characterized as "skeptical" (undertrusting).

This is an elegant theory that has some support. Moray et al. (2000) found that in an adaptive automation experiment with varying levels of automation, participants converged on the optimal level of trust—in the sense that subjective trust matched the objective reliability of the automation—relative to higher or lower trust. Moray and Inagaki's (2001) theory depends for its validation on one's being able to specify precisely what the optimal sampling rate of automation should be. This is easy when the failure rate of the automation is known. But monitoring automation is rarely the only task for the operator, and therefore the sampling rate also depends on the operator's being able to specify the required sampling frequency of all the other tasks in which he or she is engaged. In general, identifying an optimal automation sampling rate may be difficult to do in complex systems with multifunction forms of automation and simultaneous manual tasks. Moreover, the theory cannot explain why operators apparently sample less frequently when the automation failure rate remains constant than when it varies at about the same mean rate.

Parasuraman et al. (1993) found that the so-called complacency effect was considerably reduced when automation reliability was variable over time compared with when it was constant (a finding recently replicated by Bagheri & Jamieson, 2004). The overall failure rate was the same in both conditions, but operators noticed more automation failures in the variable-reliability condition than in the constant-reliability condition, perhaps because they were more skeptical.

One explanation for undermonitoring of automation that complements the trust theory is based on attention. Common sense tells us that there is nothing to be gained by attending to very reliable automation with low downside risk—for example, our home heating systems—until after failure is evident. Attending to imperfect automation is also diminished when the operator is engaged in other tasks that require focal attention. Evidence in support of this view comes from studies of eye movements. Metzger and Parasuraman (2001), for example, asked experienced air traffic controllers to monitor a radar display for separation conflicts while simultaneously accepting and handing off aircraft to and from their sector, managing electronic flight strips, and using data linking to communicate with pilots. They were assisted by a "conflict probe" aid that predicted the future (up to 8 minutes) courses of pairs of aircraft in the sector. The automation was highly reliable and reduced the time that controllers took to call out the conflict (see also Metzger & Parasuraman, 2005).

In one scenario, however, the automation did not point out the conflict because it did not have access to the pilot's intent to change course. Controllers were either considerably delayed or missed the conflict entirely. Eye movement analysis showed that those controllers who did not detect the conflict in this case had fewer fixations of the radar display compared with when they had been given the same conflict scenario without the conflict probe aid. This finding is consistent with the view that overreliance on automation is associated with reduced attention allocation, compared with manual conditions.

HUMAN PERFORMANCE RESEARCH RELATED TO AUTOMATION

Trust as a Design and Performance Issue in Human-Automation Interaction

Until recently system design engineers seldom—if ever—used the term *trust* when discussing automation, but today it is seen as a key concern (Lee & See, 2004; Parasuraman & Riley, 1997).

For different people and contexts the term can have different meanings (Sheridan, 1988). Trust can be both a cause and an effect. It is a cause in the sense that a human's use of the automation depends on his or her trust of it. As an effect, it can have several connotations:

1. judged reliability of the automation, in the usual sense of repeated, consistent functioning;
2. perceived robustness—that is, demonstrated or promised ability to perform under a variety of circumstances; sense of familiarity, whereby the system employs procedures, terms, and cultural norms that are familiar; perceived understandability, in the sense that the human supervisor can form a mental model and predict future system behavior; usefulness of the system to the trusting person; or dependence of the trusting person on the system.

Lee and See (2004) pointed out that user or system vulnerability to automation error is a critical component of the definition of trust; something important must be riding on the decision to rely on the automation. This is an important point because in numerous recent "automation" studies, the experimental participants have nothing at stake if the automation fails. In prior work, Riley (1996) demonstrated that perceived risk has a significant influence on automation reliance. His concern was that some researchers may be losing sight of the fact that automation failures in the real world can have real consequences, that human biases regarding errors of omission and commission are related to the nature of the real potential outcomes, that these biases will likely apply to automation use decisions, and therefore that studies that do not incorporate the element of operator vulnerability to automation error may yield potentially misleading results (see also Parasuraman & Riley, 1997).

Lee and Moray (1992) studied trust during the supervisory control of a simulated pasteurizing plant in which operators could select either manual or automatic control in the face of random failures. They proposed a quantitative model of dynamic changes in operators' trust and their decisions to use the automation, in which trust depended on the past level of trust, failure probability, and subjectively perceived capability in manual control. Shifts between automatic and manual control modes could be predicted by the ratio between trust in the machine and self-confidence in one's own manual performance (Lee & Moray, 1992, 1994; Muir, 1988).

An interesting follow-up study was carried out by Lewandowsky, Mundy, and Tan (2000) in which people shared control either with a human or with automation but were told in all cases that they were sharing with a human. They were more tolerant of system errors when they believed it was a human.

Engendering trust is often a desirable feature of a system, something the designer strives for—but not always. Too much trust (usually naive trust) can be just as bad as too little. Parasuraman and Riley (1997) made a compelling case that system designers should be concerned about misuse, disuse, and abuse of automation based on distrust and overtrust as well as on workload and other factors.

"Etiquette" in the Design of Human-Automation Interactive Software

In their recent review of the literature on user trust in automation, Lee and See (2004) proposed that trust is related to emotions and attitudes that people have regarding automation. As Nass and colleagues (Nass, Moon, Fogg, Reeves, & Dryer, 1995; Reeves

& Nass, 1996) showed, people often respond socially to computers in ways similar to how they interact socially with other people. Computers are becoming more intelligent—more like people, some might say. Because machines and the people who use them need to communicate and cooperate, and because people are accustomed to communicating with other people, it has become clear that the same social mores that apply to people might also apply to intelligent machines.

Individuals are typically most attracted to others who appear to have personalities similar to their own. This phenomenon, which psychologists call the *social attraction hypothesis,* also predicts user acceptance of computer software (Nass et al., 1995). These considerations suggest that etiquette, or adherence to an accepted but frequently implicit code of behavior between individuals in any social setting, may also play a key role in human-computer relations.

What is commonly considered to be good etiquette is a set of behavioral practices that make interaction between mature people acceptable and efficient (Grice, 1975). If this were not so, etiquette might have died out long ago. Grice posed four maxims for cooperation in conversation (etiquette):

1. Maxim of quantity: Say what serves the present purpose but not more.
2. Maxim of quality: Say what you know to be true based on sufficient evidence.
3. Maxim of relation: Be relevant, to advance the current conversation.
4. Maxim of manner: Avoid obscurity of expression, wordiness, ambiguity, and disorder.

Grice's axioms were considered by C. A. Miller (2004) in the context of designing adaptive user interfaces. His proposed rules, slightly abbreviated, are as follows:

1. Make many conversational moves for every error made.
2. Make it very easy to override and correct any errors.
3. Know when you are wrong, mostly by letting the human tell you.
4. Don't make the same mistake twice.
5. Don't show off. Just because you can do something does not mean you should.
6. Talk explicitly about what you are doing and why. (Your human counterparts spend a lot of time in such meta-communication.)
7. Use multiple modalities and information channels redundantly.
8. Don't assume every user is the same; be sensitive and adapt to individual, cultural, social, and contextual differences.
9. Be aware of what your user knows, especially what you just conveyed (i.e., don't repeat yourself).
10. Be cute only to the extent that it furthers your conversational goals.

The work of C. A. Miller (2004) suggests that automation that follows such agreed-upon axioms of etiquette is more likely to be accepted and liked by human operators. But is there any objective evidence that designing etiquette into automation can improve system performance? In other words, does etiquette matter as far as the bottom line is concerned—productivity, efficiency, safety—or is it simply something that makes the operator feel good, an adjunct to the principle of human-centered automation?

A recent experiment by Parasuraman and Miller (2004) suggests that etiquette can influence efficiency and therefore possibly safety. These authors examined whether etiquette affects human decisions to use automation effectively in high-workload situations and, if so, whether in the same way as automation reliability influences user trust and usage. Unreliable or imperfect automation is generally correlated with decreased trust and decreased reliance on automation, but various other factors moderate this phenomenon (Lee & See, 2004). To what extent is etiquette one such factor?

Parasuraman and Miller investigated whether good automation etiquette could compensate for poor reliability and result in increased use of automation, and, conversely, whether poor etiquette could negate the benefits of high reliability and result in decreased trust and automation usage. Such phenomena are not uncommon in human-human relationships. To test the hypothesis, 16 participants (general aviation pilots and nonpilots) were examined on a flight simulation task, the Multi-Attribute Task (MAT) Battery (Comstock & Arnegard, 1992), which incorporates a primary flight (tracking) task, a fuel management task, and an engine health monitoring task. Participants performed the first two tasks manually at all times so as to simulate a high-workload environment.

Intelligent automation support modeled after the Engine Indicator and Crew Alerting System (EICAS) that is typically installed in modern automated aircraft was provided for the engine systems monitoring task. In this task, participants had to check particular engine parameters for malfunctions and make an appropriate diagnosis. They were asked to query the EICAS system while using the automation support tool, which provided advice on possible malfunctions. An example advisory message would be, "The Engine Pressure Ratio (EPR) is approaching Yellow Zone. Please check. Also, cross-check Exhaust Gas Temperature (EGT). There is a possible flame-out of Engine 1."

Parasuraman and Miller (2004) defined good automation etiquette as a communication style that was "non-interruptive" and/or "patient." Conversely, poor automation etiquette occurred when the automation communicated in an "interruptive" or "impatient" manner. For example, in the interruptive (poor etiquette) case, the automation provided advice without warning and came on when the user was already querying EICAS and was engaged in fault diagnosis (i.e., already engaged in the behavior the system was recommending). In the impatient case, the automation urged the next query before the user was finished with the current query. (The experimenters confirmed their definitions of good and poor etiquette by questioning the participants at the end of the experiment.)

The good and poor etiquette interfaces were combined with two levels of automation reliability, following a previous study by Parasuraman et al. (1993). In the high-reliability condition, the EICAS provided correct advice on 8 of 10 engine malfunctions (80%), whereas it was correct on only 6 of 10 malfunctions (60%) in the low-reliability condition. The results were, as expected, that user diagnostic performance was better when the reliability of the automation was high (80%) than when it was low (60%). However, and less obviously, good automation etiquette significantly enhanced diagnostic performance both when automation reliability was low and when it was high. Third, and perhaps most interestingly, the effects of automation etiquette were powerful enough to overcome low automation reliability: Performance in the low-reliability/good-etiquette condition was almost as good as (and not significantly different from) that in the high-reliability/poor-etiquette condition. These performance findings were paralleled by similar effects on user ratings of trust in the automation.

Finally, the effects of poor etiquette on human-system performance were not attributed solely to distraction or interruption. All messages, whether associated with good or poor etiquette, were presented visually in the EICAS communications window and not, for example, using speech synthesis, which could have been distracting. Furthermore, a control group of participants was run to test the hypothesis that any interruption might be expected to degrade user performance. These participants received nonspecific interruptions—for example, "Maintaining primary flight performance is important, but do not forget to check engine parameters for possible malfunction." These neutral interruptions had no effect compared with the good and bad etiquette messages.

These results provide strong evidence for the influence of automation etiquette on both user performance and trust in using an intelligent fault management system to diagnose engine malfunctions, at least in a high-workload setting when users are busy doing other tasks. The results also clearly show that reliability per se may not be sufficient to promote overall human-machine system efficiency: Both user diagnostic performance and trust were lowered by poor automation etiquette even when the reliability of the advice provided by the automation was relatively high.

The results also suggest the intriguing notion that good automation etiquette can compensate for the performance costs associated with low automation reliability. Some may find this result disturbing, for it suggests that developing robust, sensitive, and accurate algorithms for automation—a challenging task under the best of circumstances—may not be necessary so long as the automation "puts on a nice face" for the user. This is unlikely, however, because this experiment also clearly showed that the best user performance (and the highest trust) was obtained in the high-reliability condition in which the automation also communicated its advice to the user with good etiquette.

Ecological Interface Design and Automation

Systems that provide the user with information concerning automation modes, system states, and future automated actions—particularly if they do so with good etiquette—can improve human—automation communication and therefore potentially enhance system performance. If the provision of such feedback requires extensive cognitive processing on the part of the user, however, any benefit may be counteracted by the increased cognitive load on the user. Consequently, there is a need to develop interfaces that provide feedback on automation states and behaviors in a manner that requires little or no cognitive effort but can be directly apprehended by a quick glance at a display that provides the appropriate avenue for rapid action. Such interfaces have collectively been characterized under the rubric of *ecological interface design* (EID; Vicente, 2002; Vicente & Rasmussen, 1992). The origin of the term *ecological* can be traced to the perceptual theories of Gibson (1979), whose concepts of *affordance* and *direct perception* are closely paralleled by EID.

The EID approach also draws on the concept of the goals-means-ends abstraction hierarchy proposed by Rasmussen (1986) and on his taxonomy of skills, rules, and knowledge. In this approach a given domain of work is decomposed into parts, beginning with the top goal, then into the means by which the goal is achieved, both in terms of abstract functions and, at the bottom of the hierarchy, in terms of the physical acts needed to achieve those functions. In theory, EID interfaces are thought to allow for direct perception of functional relationships between system components, without the need for extensive cognitive processing, and to enable rapid action. In essence, the EID interface is thought to replace cognition with perception, thereby facilitating action. Accordingly, human-automation interaction might also be facilitated if the interfaces used to provide feedback to the user on automation states are designed in accordance with EID principles.

Two studies reported by Furukawa, Parasuraman, and Inagaki (2003) are described for the purpose of illustrating the efficacy of EID interfaces for human-automation interaction. In the first study (see also Molloy & Parasuraman, 1994), an integrated display with an emergent perceptual feature was used to examine pilot fault management performance in a flight simulation task. As discussed previously, when operators are busy with several tasks, they may fail to monitor the status of an automated task, so that when occasional failures occur, operators are less likely to detect the failure (or are slow in responding) compared with when they perform the task manually (Parasuraman et al., 1993). This automation complacency effect reflects a policy of allocating attention away from the automated task to the other tasks that the operator has to perform simultaneously (Metzger & Parasuraman, 2005; Moray & Inagaki, 2001; Parasuraman et al., 1993). Hence, one way to mitigate the effect would be to use an integrated display (Bennett & Flach, 1992) based on EID principles to display system states, both under normal conditions and when malfunctions occur.

In the first study reported by Furukawa et al. (2003), 12 general aviation pilots performed a flight simulation task involving the following subtasks: two-dimensional compensatory tracking (a resource management task requiring balancing of the fuel tanks of the aircraft) and an automated engine systems failure detection task. The interface used to display the engine systems task consisted of either an integrated or a nonintegrated display. The pilot's task was to detect deviations in different engine parameters (EPR, N1, EGT, etc.) and take appropriate corrective actions.

The nonintegrated display was based on a traditional EICAS, which was displayed with circular gauges showing engine parameters in an analog display. The integrated display was based on the Engine Monitoring and Crew Alerting System (EMACS), which was depicted with a deviation bar graph. The tips of the four bars displayed an emergent perceptual feature—an implied contour line when the state of the engine was normal—which provided a ready indication of system normality or malfunction without the expenditure of much cognitive effort. In Rasmussen's (1986) terms, the line provided the abstract information obtained from the state parameters during normal operation. If one of the engines failed, the emergent line was distorted, enabling the pilots to detect it easily.

There were no differences in pilot performance between the EICAS and EMACS displays under normal conditions. However, when the automation failed to detect and diagnose malfunctions, pilot monitoring performance (detection and diagnosis rates) was significantly better for the integrated EMACS display than for the nonintegrated EICAS display. Moreover, the automation complacency effect—the reduction in performance for automated compared with manual monitoring—was eliminated with the integrated display. Analysis of eye movements indicated that dwell times on the display were significantly shorter with the EMACS than with the EICAS display.

These findings indicate that automation complacency results from a policy in which attention is allocated away from the automated task to the other manual tasks that the operator has to perform. As a result, if abnormalities in the automated control of a task occur, malfunctions may not be rapidly detected and diagnosed. However, if these are shown to the operator using an integrated display based on EID, fewer attentional resources are needed to detect and diagnose the malfunction. This was supported by the eye movement analysis: Pilots had shorter visual fixations (dwells) with the EMACS display, yet they performed better than with the EICAS display, suggesting that display integration facilitated efficient information extraction. Furthermore, EID eliminated the automation complacency effect.

In a second experiment, Furukawa et al. (2003) used a modification of the DURESS process control simulation (Vicente & Rasmussen, 1992) to which automated controllers with three different modes were added. As a result, participants had considerable difficulty in recognizing automation states and changes using a standard display of the process control simulation. EID was used to design a new display that supported operators in their ability to prevent conflicts with the automated controllers by explicitly representing the intentions (i.e., goals and means) of the automation on the display. Performance in 16 participants was compared with the intention-represented EID and a standard EID display. Participants were shown one of the displays for 8–10 seconds and then were asked to report key parameters in the current state of the system as well as its future state. They also had to indicate the control operations that needed to be taken and the target states of the operations. The scenarios included situations that had been experienced before as well as unexpected events.

The results showed that with such a mode-rich, high-capability automated system, operators found it difficult to recognize the goals, the means, and their interrelations (means-end relations) of the automated controllers using the standard display. Some participants failed to supervise particular tasks performed by the automated controllers, perhaps because they thought that controllers had complete responsibility for the tasks, and even though they clearly demonstrated their comprehension of the situation. However, performance was enhanced when the intention-represented EID display was used. With this display, the automated controllers revealed their goals and means to participants, like human supervisors or coworkers usually do. As a result, participants could recognize the means-end relations through observing the behaviors of the automated controllers. Consequently, they took appropriate actions that were not in conflict with those of the automatic controllers, thereby enhancing performance.

In general, these results indicate that EID displays might be particularly helpful for human-automation interaction under nonroutine conditions. Under routine operating conditions that have been experienced many times before, human operators can effectively monitor and control a complex system using standard operating procedures. Although this is efficient for normal operations, it may not be appropriate for abnormal or unanticipated conditions. When unexpected events occur, a human-automation display interface is required that allows for quick comprehension of system state and rapid action.

QUANTITATIVE MODELS

Quantitative Models for Supervisory Functions

As discussed previously, human-automation interaction can require several different functions served by the human supervisor and/or the automation, as shown by the boxes at the top of Figure 2.1. Quantitative models for each of these supervisor functions are available in the literature (see Table 2.3).

Most of these models were developed by engineers for engineering applications that had little to do with human interaction (and mostly in conjunction with World War II). Since 1950, many such engineering theories have been adapted to the analysis and modeling of human behavior. For example, following Shannon's exposition of information theory (1947) applied to telephone communication systems, psychologists modeled a variety of sensory-motor skill tasks (Hick, 1952) using an information transmission model (bits per second), and G. A. Miller (1956) applied information transmission (in bits per selection) to performance in immediate memory tasks.

Following the development of signal detection theory for sonar and radar detection applications, Green and Swets (1966) and Swets (1996) applied that theory to sensory psychophysics. And following the development of classical automatic control theory for gun control (James, Nichols, & Phillips, 1947; Craik, 1947), McRuer and Krendel

TABLE 2.3: Theoretical Models Relevant to the Five Supervisory Functions

Function	*Appropriate Type of Quantitative Model*
Plan	
Model physical system	System analysis/differential equations
Satisfice trade-offs among objective	Multiattribute utility theory
Formulate control strategy	Control theory, optimization theory
Teach	Information and coding theories
Monitor	
Allocate attention	Sampling theory, Markov network analysis
Estimate state	Estimation theory
Detect failure	Signal detection, Bayesian analysis, pattern recognition theories
Intervene	Decision theory
Learn	Human and machine learning theories

(1959) and McRuer and Jex (1967) applied the control-loop model to human pilot control of aircraft in turbulence. Kleinman, Baron, and Levison (1970) applied the more recent optimal control theory to pilot instrument scanning and control of an aircraft during approach.

These and other adaptations of engineering theories to simple sensorimotor skills have been reviewed by various authors (Rouse, 1980; Sheridan & Ferrell, 1974), but none of the aforementioned modeling efforts really dealt with human-automation interaction. Many classical papers (including the classical control papers mentioned earlier) were recently reprinted in Moray (2005).

Understanding the physical interactions of the processes to be controlled, setting of goals and constraints, and devising a control strategy (all part of planning); communicating the model, the goals, constraints, and control strategy to a computer (teaching); allocating attention and estimating system state (monitoring); preparing for emergency action (intervening); and learning from experience (learning) are all relatively high level cognitive tasks—and we really do not have an acceptably robust quantitative model for even one of these functions, let alone all of them working together. So, in spite of the seeming relevance of the above-listed engineering models to the various human supervisory functions, the challenge remains to make the relevance explicit and to devise tractable predictive models.

Comparing Performance Attributes of Humans and Automated Systems

One type of model that is relatively easy to realize is that in which certain characteristics of an automatic system can be specified and comparisons made with either a human or another automatic system in terms of those characteristics. For example, Sheridan and Parasuraman (2000) offered a simple analytical criterion for deciding whether a human or automatic system is better in a failure detection task. The method is based on expected-value decision theory in much the same way as is signal detection. It requires specification of the probabilities of misses (false negatives) and false alarms (false positives) for both the human and the automation being considered, as well as factors independent of the choice—namely, costs and benefits of incorrect and correct decisions and the prior probability of failure. The method can also serve as a basis for comparing different modes of automation. The authors discuss some limiting cases of application as well as some decision criteria other than expected value.

ADAPTIVE AUTOMATION

Thus far we have discussed the effects on system performance of various forms of automation, whether defined by different levels of autonomy or by different stages of information processing to which machine aiding can be applied. One goal of much of this research is to influence system design by recommending a particular level or type of automation or, as in the model by Sheridan and Parasuraman (2000), deciding whether to automate at all. As such, this work assumes that once the final design has been identified, the automation has a fixed level or type consistent with these design features. This approach, in which the characteristics of automation are set at the design stage and then executed in the same way during operational use, has been referred to as *static* automation (Parasuraman, Bahri, Deaton, Morrison, & Barnes, 1992). In contrast, in *adaptive* automation, the level and/or type of automation is not fixed but may change during system operations.

Adaptive automation is akin to *dynamic* function allocation, in which the division of labor between human and machine agents is not fixed but dynamic, flexible, and context dependent (Hancock, Chignell, & Lowenthal, 1985; Kaber & Endsley, 2004; Parasuraman, 1993; Parasuraman et al., 1992; Rouse, 1988; Scerbo, 1996, 2001; Inagaki, 2003). For example, if performance in a higher level of automation is getting worse, the automation may change to a lower level and/or turn over more or even all control to the human. If, on the other hand, high human workload is detected and/or the human is not responding appropriately, automation may go to a higher level so as to become less dependent on the human.

The thorny human factors issue of allocation of function has typically been based on stereotypical characteristics of human and computer capabilities, an approach that has met with marginal success (Dialogs on Function Allocation, 2000; Jordan, 1963; Sheridan, 1998). Hancock et al. (1985) suggested that the issue and its attendant problems may be bypassed if function allocation is viewed as dynamic rather than static. In this view, allocation of function is not just a design activity but something that occurs during system operations. Accordingly, there has been interest in investigating the effects of different forms of adaptive automation on human performance in various simulated tasks.

Techniques for Adaptive Automation

Although the adaptive automation concept is not new (Rouse, 1976), technologies for its effective implementation have been explored only recently. Key technology development efforts include the Air Force's Pilot's Associate (Hammer & Small, 1995), the Navy's Adaptive Function Allocation for Intelligent Cockpits (Morrison & Gluckman, 1994; Parasuraman et al., 1992), and the Army's Rotocraft Pilot's Associate programs (C. A. Miller & Hannen, 1999). At the same time, several researchers have investigated human performance in relation to adaptive automation using simulations of flight, air traffic control, driving tasks, and process control (Inagaki, 2003; Moray et al., 2000; Parasuraman, 1993; Scerbo, 1996).

A fundamental issue in all these areas concerns the means by which adaptive automation is invoked. In an early review, Parasuraman et al. (1992) identified five main categories of techniques for implementing adaptive automation: critical events, operator performance measurement, operator physiological assessment, modeling, and hybrid methods combining one or more of these techniques.

In the critical events method, automation is invoked if certain external events occur, but not otherwise. For example, Barnes and Grossman (1985) described an aircraft air defense system in which the beginning of a "pop-up" weapon delivery sequence led to the automation of all defensive measures of the aircraft. However, if this critical event did not occur, the automation was not invoked. There are many other applications in which the occurrence of critical events triggers automation of some kind.

The critical events method is flexible in that it can be tied to current tactics and doctrine during mission planning. A disadvantage of the method is its possible insensitivity to actual system and human operator performance. For example, this method will invoke automation irrespective of whether or not the pilot requires it when the critical event occurs.

Measurement of operator performance and physiological measurement may be used to overcome this limitation. Here the goal is to change the level or type of automation based on an assessment of operator states. Various operator states (e.g., mental workload, fatigue, or—more ambitiously—operator intentions) may be inferred on the basis of performance or other measures and the resulting value input into the adaptive logic of an expert system or neural network. Kaber and Riley (1999), for example, used a secondary-task measurement technique to assess operator workload in a target acquisition task (see also Kaber & Endsley, 2004). They found that adaptive computer aiding based on the secondary-task measure enhanced performance on the primary task.

Operator physiological assessment offers another potential input for adaptive systems (Byrne & Parasuraman, 1996; Parasuraman et al., 1992). For example, physiological measurements may indicate that a human operator is dangerously fatigued or experiencing extremely high workload. An adaptive system could use these measurements to provide computer support or advice to the operator that would mitigate the potential danger. Technology is available to measure a number of physiological signals from the operator, from autonomic measures such as changes in heart rate to central nervous system measures such as the EEG and event-related potentials (ERPs), as well as measures such as eye scanning and fixations (see Scerbo et al., 2001, for a review).

The main advantage of (at least some) physiological measures is their high bandwidth compared with most performance measures (with the exception of performance in manual tracking). There is now a substantial literature indicating that several psychophysiological measures can be used for real-time assessment of mental workload (Scerbo et al., 2001). Prinzel, Freeman, Scerbo, Mikulka, and Pope (2000) also specifically demonstrated the feasibility of an adaptive system based on EEG measures.

Wilson and Russell (2003) used an artificial neural network (ANN) with multiple physiological measures to identify low- and high-workload periods in real time during performance of the MAT flight simulation. When the ANN detected a high workload state, two of the MAT subtasks were automated. Wilson and Russell (2003) found that implementing adaptive automation with the ANN method led to improved performance on the MAT subtasks compared with manual performance. However, one limitation of this study is that no comparison to static automation was made, so that the specific benefit of adaptive automation was not identified.

The four major techniques for implementing adaptive automation have complementary merits. The critical event technique has the advantage that it can be tied to mission planning but the disadvantage that it does not take into account operator requirements. Measurement of operator performance or physiological state has the advantage of being potentially responsive to unpredictable changes in human operator cognitive states. Physiological measures can be designed to be relatively unobtrusive, but their sensitivity and validity need to be established in each application domain.

Moreover, all operator assessment methods are only as good as the sensitivity and diagnosticity of the measurement technology. Modeling techniques have the advantage that they can be implemented offline and easily incorporated into rule-based expert systems (Parasuraman et al., 1992). However, this method requires a valid model, and many models may be required to deal with all aspects of human operator performance in complex task environments. Furthermore, different models might give contrary decisions at a particular moment.

For these reasons, therefore, hybrid methods that attempt to optimize the relative benefits and disadvantages of each of these techniques may offer the best general approach to implementing adaptive automation. Many investigators are attempting to develop such hybrid approaches to adaptive automation (e.g., St. John, Kobus, Morrison, & Schmorrow, 2004), but nothing practical has been produced to date.

Mitigation of Driving Distraction by Adaptive Information Systems

An active current research area in adaptive automation is to make information technology in highway vehicles adaptive to traffic and other workload demands (Lee, Caven, Haake, & Brown, 2001; Lee, McGehee, Brown, & Reyes, 2002). There is concern that drivers are being distracted from the driving task by use of cell phones, radio and other entertainment systems, navigation systems, personal digital assistants (PDAs), and, possibly in the future, e-mail, faxes, stock trading, and all sorts of real-time interactions with personal or corporate computer files. The highway vehicle can easily become the office. The problem is one of maintaining safety.

If distraction could be measured, it could be mitigated by various interventions. Possible measurements of distraction include eye fixations off the road (measured by eye-tracking devices), high traffic density (measured by out-the-window video image analysis), control actions (measured by spectral analysis of steering wheel adjustments, movement of the foot off the accelerator pedal, and use of turn signals), weather (measured by precipitation sensors, road surface sensors, and use of windshield wipers), and excessive speed. Interventions could be warning signals (e.g., auditory warnings in the cell phone causing the other party as well as the driver to be aware of the traffic density or distraction source), disabling of certain functions (cutting off the cell phone, navigation system, radio, or whatever), and modifications to other systems (making the intelligent cruise control follow at a greater distance from the lead vehicle, or providing a forward collision warning sooner than would otherwise occur).

The National Highway Traffic Safety Administration has sponsored several national forums on the topic of driver distraction (Llaneras, 2000). The detailed research questions relating to distraction mitigation by automation are too numerous to list here.

Adaptive Automation in Air Traffic Control

Unlike the aircraft flight deck, the air traffic control (ATC) center typically has not been highly automated. Because of increased air traffic, however, and given the need for greater efficiency, there have been several proposals to modernize the ATC system with increased automation. A National Research Council committee examined this issue of ATC modernization and made a number of recommendations for enhancing the current system with judicious use of automation of certain controller functions (Wickens, Mavor, Parasuraman, & McGee, 1998). Some of these recommendations appear to have been followed by the FAA, which has begun the implementation of a number of new automated systems designed to aid controllers in the safe and efficient separation of air traffic.

One example of new ATC automation is the Center TRACON Automation System (CTAS), which consists of a suite of software tools for aiding controllers. Hilburn, Jorna, Byrne, and Parasuraman (1997) used one of these tools to examine the effects of adaptive automation on the performance of air traffic controllers. Previous research found that automation that is introduced in an attempt to reduce the operator's mental workload often has the opposite effect: increasing mental workload during high task demand periods, reducing it unnecessarily during periods of low demand, or both (E. L. Wiener & Curry, 1980). This phenomenon, one of the "ironies of automation" first referred to by Bainbridge (1983) and which E. L. Wiener (1988) called "clumsy" automation, could be mitigated if the onset of automation were linked to task demand.

Hilburn et al. tested experienced controllers on a high-fidelity simulation of the Brussels-Maastricht ATC sector. They varied the traffic load and complexity to create low and high task demand periods during a simulated work shift. The controllers were provided with a decision aid for determining the optimal descent trajectories of aircraft at the start of initial approach—the Descent Advisor (DA) of CTAS. The DA also gave information on potential aircraft-to-aircraft or aircraft-airspace conflicts and offered possible resolutions to the conflict.

Hilburn et al. used two conditions: in the static automation condition, the DA was provided throughout the work period, irrespective of traffic load, whereas in the adaptive automation condition, the DA was provided only at high traffic load and not during periods of low traffic load. These researchers found significant benefits for controller workload (as assessed using pupilometric and heart rate variability measures) when the DA was provided adaptively during high traffic loads, compared with when it was available throughout (static automation) or only at low traffic loads. In addition, the adaptive condition was associated with significantly better system performance, as indicated by a lower difference between estimated and actual time of arrival at the airport and higher throughput.

Interfaces for Adaptable Automation: Delegation

Although there is a fairly large body of empirical evidence pointing to the system performance benefits of adaptive automation (Inagaki, 2003; Parasuraman, 2000; Scerbo, 2001), it is far from clear how well fully adaptive systems will perform in practice. In adaptive systems, the decision to invoke automation or to return an automated task to the human operator is made by the system in real time using any of the previously described methods. This immediately raises the issue of user acceptance of such a system. Human operators may be unwilling to accede to the "authority" of a computer system that mandates when and what type of automation is or is not to be used.

Also of concern is the potential for system unpredictability and its consequences for operator performance. Billings and Woods (1994), for example, warned that truly adaptive systems may be problematic because their behavior may not be predictable to the user. To the extent that automation can hinder the operator's situation awareness by taking him or her out of the loop, unpredictably invoked automation by an adaptive system may further impair the user's understanding of the situation. However, if the automation were explicitly invoked or changed in mode by the user, then presumably the unpredictability would be lessened.

But involving the human operator in making decisions about when and what to automate can increase mental workload. Further, in a team situation, one team member may reconfigure the system and the other team member(s) not know about it (Moray, 1992). Or, if responsibility is not clearly assigned, one team member may mistakenly assume that another team member will take some necessary action.

Thus, there is a trade-off between increased unpredictability and increased workload in systems in which automation is invoked by the system or by the user, respectively. Opperman (1994) characterized these alternatives as *adaptive* and *adaptable* approaches to system design (see also Scerbo, 2001). The combined human-machine system adapts to various contexts in both cases. However, in adaptive systems, automation determines and executes the necessary adaptations, whereas in adaptable systems, the operator is in charge of the desired adaptations. The distinction is primarily one of authority.

In an adaptable system, the human always retains the authority to invoke or change the automation, whereas in an adaptive system this authority is shared. In adaptable systems, therefore, the human operator is like the supervisor of a human team who delegates tasks to team members—or, in this case, to automation. The challenge for developing *delegation interfaces* to a system is that the operator should be able to make decisions regarding the use of automation in a way that does not create such high workload that any potential benefits of delegation are lost.

Delegation interfaces may allow adaptable automation to be implemented at a flexible and appropriate balance point in this trade-off space (C. A. Miller & Parasuraman, in press). Humans should be able to delegate tasks to automation at times of their own choosing and receive feedback on their performance. Delegation in this sense is similar to that which occurs in successful human teams—for example, self-organizing of roles on aircraft carriers (Rochlin, LaPorte, & Roberts, 1987; LaPorte, 1996). It represents a real-time approach to supervisory control (Sheridan, 1976, 1992b). The human operator sets an objective, provides instructions (at a greater or lesser level of detail), and then delegates or authorizes the automation to determine the best method by which to proceed within the instructions toward the goal. Delegation should provide a highly flexible method for the human supervisor to declare goals and provide instructions and thereby choose how much or how little autonomy to impart to automation on a moment-by-moment basis.

An example of such a delegation interface is the Playbook™—so named because it is based on the metaphor of a sports team's book of approved plays and the selection from among those plays by the team leader—for example, the quarterback in American football—and their execution by the team members—that is, the other players

(C. A. Miller & Parasuraman, in press; C. A. Miller, Pelican, & Goldman, 2000). The Playbook interface facilitates the teaching of automation, an idea first proposed by Sheridan (1976, 1992a) when he suggested the supervisory control concept. This is both a shared knowledge structure of tasks and their relationships within which task performance can be discussed by human and automation and a language for the communication of instructions.

The Playbook uses a hierarchical task model to provide a common language with which a human supervisor may communicate goals and intents, and a Hierarchical Task Network planning system (Erol, Hendler, & Nau, 1994) to understand, reason over, and either critique or complete partial plans provided by the human. This form of interface permits the operator to delegate tasks to automation at a wide variety of functional levels of abstraction by provision of goals and of full or partial plans. Finally, the Playbook streamlines the process of delegation by the human operator by providing a compiled set of plans, or "plays," with short, easily commanded labels that can be further modified as needed. This is the critical aspect of the concept that allows this form of adaptable automation not to increase the workload associated with delegation, much as a sports team has an approved set of plays that facilitate task delegation by the team leader.

There are two sources of evidence concerning the efficacy of the delegation approach to adaptable automation. First, a Playbook prototype for a mission planning tool for commanding unmanned combat air vehicles (UCAVs) has been developed as a proof of concept (C. A. Miller, Goldman, Funk, Wu, & Pate, 2004). Second, initial experimental studies of the effects of delegation interfaces on human performance have been carried out (Parasuraman, Galster, Squire, Furukawa, & Miller, 2005). These studies examined the use of a simple delegation interface on system performance during simulated human-robot teaming using the RoboFlag simulation environment. RoboFlag provided the operator the ability to command simulated robots, individually or in groups, at several levels of detail: by providing designated endpoints for robot travel, by commanding higher-level behaviors (or modes or plays) such as "Patrol Border" or "Circle Defense," or by even higher "super-plays" such as "Go on Offense." The results showed that the multilevel tasking provided by the delegation interface allowed effective user supervision of robots, as evidenced by the number of missions successfully completed and the time for mission execution. However, additional studies are needed in which more complex versions of delegation interfaces are evaluated.

CURRENT AND FUTURE AUTOMATION TECHNICAL CHALLENGES

Future automation technology will be characterized by being smaller and smarter. Miniaturization of sensors, actuators, and computers (nanotechnology) goes hand in hand with lower power consumption. This will mean that the human user will command, directly or indirectly, many more automatic control loops than today. However, most of these will not be apparent to the human; they will be hidden from view, much as are the 20–40 computer chips in today's automobiles. More and more of the analysis and decision making will be provided by so-called computerized intelligent agents (Weiss, 1999), but there are serious challenges in getting humans and automation to cooperate, as will be discussed next.

Challenges in Various Application Sectors

Sectors of the economy where robot automation has been prevalent will see a continuing trend toward more, smaller, and more capable robots in manufacturing plants, in hostile environments such as undersea and in space, on farms, in homes, and in defense systems.

A visitor to a modern automobile final assembly plant will find numerous robots fitting parts together (including glass and cloth or rubber components that only a few years ago required skilled human labor) as well as welding and painting. Modern chip-making and electronics assembly are done largely by robots, which have the advantage of being faster, cleaner, and more precise than humans. In chemical processing and genetic engineering, robotic manipulators combine with automatic measurement and flow control to perform batch processing, often on a scale (of complexity, not size) that is far greater than what a human operator can handle, or far more delicate, and in most cases much faster. Even in China, where labor is cheap by Western standards, engineers are installing factory automation on a large scale. In all such cases the human operator is relegated to the role of supervisor, as described previously.

The undersea robot (controlled remotely by a human) is gradually replacing the human diver because it is able to go to the greatest ocean depths and to work continuously. Robots are now performing oil, gas, and mineral exploration and mining operations in the ocean, where operations must gradually go deeper as land-based and shallow-ocean reserves run out. Space robots have repeatedly proven themselves well suited to planetary exploration and more recently have been favored over human astronauts for performing delicate repairs on the Hubble space telescope and other valuable scientific instruments.

Most undersea and space robots do not perform repetitive tasks as do factory robots; each new task is different. They are more appropriately called *teleoperators* (if they are remotely controlled by humans continuously) or *telerobots* (if they are truly autonomous at least for short periods and generally under supervisory control by their remote human operators; see Sheridan, 1992a).

As world populations increase, deforestation continues, and arable lands slowly become desert or are flooded by rising ocean levels, the demand for food will increase. A demand for more automation is bound to occur to make farming more efficient. Robotic tractors and harvesters are gradually replacing manned vehicles on farms. Agricultural robotics will include robots to farm the ocean, which accounts for roughly 95% of the earth's biosphere (by volume) and is a good source of both animal and plant protein. However, despite all the progress in robotic sensing, mobility, and manipulation, the robot brains are still not capable of being independent of the human supervisor.

For the most part there has been no need to make robots look like people (anthropomorphic, or humanoid), except for entertainment. However, roboticists have shown increasing interest in research to provide robots with mobility, gestures, and facial expressions that resemble those of humans. This is in part a means to study human behavior (the idea being that imitation begets understanding) and partly to improve communication between robots and people; presumably a person can work better with a machine that acts like a person, though this premise needs to be more fully explored (Lewandowsky et al., 2000). Such “social robots” are also being designed for use with special populations of people, such as the frail elderly, the disabled, or autistic persons.

Until now the World Wide Web has been used primarily for communicating text, graphics, and software. A continuing application has been information searches, such as are provided by Google and other search engines. These are becoming ever more powerful for eliciting from a human what he or she wishes to find and for presenting that information in a manner tailored to the individual user. New applications of the Web are to control machines from a distance—for example, to start food cooking on a stove, to start the robotic lawn mower or vacuum cleaner, to control a robot to do surveillance at a home or company or government property, or to remotely manipulate or position some object for examination through a television system.

One particular automation project of special interest to human factors engineers is the effort to reduce the crews of ships, both navy and merchant marine. Traditionally, personnel on ships have been specialized: Some tended engines, some scrubbed decks, some managed cargo, and some cooked. On navy ships there were specialists for shooting guns and fighting fires. Ships of the future will require that crew members serve multiple functions and, more than anything, that they be supervisors of automation. Control will be mostly centralized in one command center. The configuration of the navy ship will be different, with very little space on open deck (practically nothing for a human to do there!). Almost all ships can now be controlled by computer; GPS sensors measure the ship’s position relative to established navigation databases of the harbors, including channels and hazards.

Cargo ships have already undergone much automation, and this trend will continue. Probably the most automated port is that at Rotterdam, where shipping containers are slid from conventional trucks onto a small number of driverless trucks, which in turn deliver them under computer control to specified points adjacent to the docked ships. Here huge robot arms pick up the containers and stack them on the appropriate ship. Unloading and stacking for subsequent dispatch is similarly automated. Human operators perform their supervisory functions from a control tower similar to that at an airport.

On naval ships there is also a general trend for fewer personnel to be available, owing to both the all-volunteer force and demographic trends. Furthermore, the multiplicity of U.S. military commitments around the world means that personnel are spread thin, and automated systems will increasingly take up the slack.

High-frequency radio technology has made many advances recently not only for cell phones and wireless computer communication but also for many other applications.

The personal digital assistant (PDA) combining radio, MP3 player, cell phone, pager, date book, and Internet access is gradually emerging in an easily wearable unit, and in the future we may no longer need these functions embodied in separate items in our homes. The first experiments are now under way with highway vehicles that communicate with one another and issue mutual warnings to vehicles converging at intersections (which can also apply the brakes should the human drivers not be paying attention).

Another example of innovative automation is radio frequency identification (RFID) chips that, when energized from a few meters away, will communicate back a radio signal. These are used on library books and retail merchandise to keep track of items going in and out of libraries and stores. RFID chips are getting tinier and cheaper (postage stamp size, and a penny or so in cost). Thus they can be pasted onto any package to maintain inventory control from the factory to the customer, with each RFID package uniquely identified by a string of bits. One can imagine all sorts of future uses to enable people to keep track of items in the home—and children. Parents in Korea are already keeping track of their children through graphic maps on their cell phones, driven by the GPS chips in their children's phones.

Automation is present in the hospital in many forms. Medical records are now computerized in many hospitals and in the future will follow the patient between the physician's office and the hospital. They will even be available to the patient, with pointers to explanatory material to educate the patient as to findings and interpretations. But this communication also poses huge problems of privacy, similar to those that have already occurred in financial transactions.

Much more automated medical testing technology will soon find its way into physicians' offices and nursing homes, and even into the home for use by patients on themselves or by family members of home health care workers. It will need to be fail-safe and easy for unsophisticated persons to use. Much-improved telecommunication between homes and hospitals will occur. This will enable home-rendered tests to be analyzed in hospitals and emergencies to be triaged so that patients will know when there is an emergency and immediate hospitalization is required and when there is no immediate cause for concern (Sheridan & Thompson, 1994).

Miniaturization also means that health sensors and alerting devices will increasingly be worn and integrated with clothing or be implanted within the body. This form of automation will have a growing market among older adults, who are making up an increasing share of our population.

Homeland security has seen much new automation to detect explosives and radioactive agents in baggage and on persons boarding aircraft. Similar new technology may soon be used on trains and to control entry to public buildings, sporting events, and so on. Improved personal identification in some form will probably be in wide use within the decade. As with health care technology, sensors to monitor personal security and alert family or police will find new markets.

Challenges of Human Operator Coordination with Automation

The new forms of automation that will pervade all aspects of life will pose new challenges for coordination between humans and automation. Woods (1996) pointed out that adding automation with the intention of assisting the human operator is like adding another team member, one who does not necessarily speak the same language and share the same cultural assumptions. This can result in what Woods has called *automation surprises* (Sarter, Woods, & Billlings, 1997) and lead the human to ask E. L. Wiener's (1982) familiar questions: What is the automation doing now? Why is it doing that? What will it do next?

Woods implies that increasing automation necessarily increases the demand for coordination. It surely does so to a greater extent than most automation designers appreciate. Yet greater coordination need not mean that the human cognitive load increases over what is required for the no-automation case. If the automation can perform assigned tasks sufficiently reliably and transparently, and if the human operator is sufficiently well trained so as to be easily able to observe what the automation is doing, to understand why it is doing that, and to predict what it will do next, the cognitive load can be diminished.

So the problem, as emphasized by Christoffersen and Woods (2002), is not one of authority or autonomy but, rather, one of cooperation and observability. And cooperation between human and machine, just as between two humans, means shared representations—where the operator's mental representation (mental model) truly corresponds to the functional and causal behavior of the machine and both correspond to the physical representations of the operator's interface (displays and controls).

Sheridan and Verplank (1978) discussed the importance in well-defined tasks—for example, in controlling a telerobot—of displays and hand controls that maintain kinematic correspondence to the posture and movement of the robot's end effector using so-called resolved motion algorithms, and thus abiding by the all-important design principles of observability and stimulus-response compatibility. Christoffersen and Woods noted that in older hardwired control rooms, multiple operators can infer what others are doing by direct visual observation, often without having to resort to symbolic voice or computer-mediated communication. Exactly how they do this, and how any team of operators uses both body and speech communication to perform joint cognitive activity—what Hutchins (1995) calls "cognition in the wild"—remains a hard problem for anthropologists and psychologists to model in any objective way. If we understood, maybe we could even endow our computer agents with a modicum of such capability.

Klein, Woods, Bradshaw, Hoffman, and Feltovich (2004) amplified these ideas by posing 10 challenges for making automation a team player in joint human-agent activity. Their point of view is a "basic compact" or tacit agreement among the joint human workers. The automation must be designed to buy into this compact, a design that includes (a) common grounding, (b) the ability to model each others' intents and actions, (c) interpredictability, (d) amenability to direction, (e) an effort to make intentions obvious, (f) observability, (g) goal negotiation, (h) planning and autonomy support, (i) attention management, and (j) cost control. As spelled out by the authors, many of these desirable goals have some overlap and have long been discussed as shortcomings of automation and questioned as being within the state of the automation design art.

Having a computer share assumptions with a human is not easy, given that every human has a lifetime of cultural assumptions that can at times (especially in decision trade-offs) bear on the task at hand. Humans have a fuzzy knowledge and rule base that is mostly beyond the capability of software engineers to encode. Having a computer model for predicting human intentions and behavior is a tall order. But these are nevertheless worthy research and engineering challenges for the long term.

SOCIAL, POLITICAL, AND ETHICAL ISSUES

Reliability and Liability

When automation fails, there are serious issues about who is liable: the operator, the firm or agency that owns the equipment and/or employs the operator, the designer, the manufacturer, the installer, the maintainer, and so on. There is a tendency to blame automation mishaps on the operator because he or she is typically present when an accident occurs. A broader view suggests that operators should have some relief from responsibility if the machine fails. Probably the wisest policy is to abstain from blaming anyone, at least until a thorough investigation has taken place. Typically it is discovered that multiple factors came together to “cause” an accident, and either the blame should be shared or no one should be blamed: the mishap should be regarded generally as a learning experience for all (Reason, 1990).

Often what the operator thinks was commanded of the automation (in the operating mode the operator thinks is operative) is not what actually was commanded. The purported father of cybernetics, Norbert Wiener, in his Pulitzer Prize-winning book (1964), asserted that as computers and automation become more complex and people become more dependent on such systems for transportation, communication, health care, and national security, there is a growing danger that the expectations of the humans will not match the logic of the machine, where the latter dictates what the automation will eventually do. Thus a degree of skepticism is critical when proponents urge that large-scale systems such as air traffic control and guided missile weapons systems be highly automated.

Virtual Reality and What Is Real

Computer graphics and display technology have made remarkable progress in recent years, enabling simulation and accompanying immersion that gives the user the feeling of actually being there. It makes human-in-the-loop simulators for aircraft, highway vehicles, ships, and minimally invasive surgery more realistic and therefore is more acceptable than it once was for training and research. Researchers have debated how and even whether such immersion is important for training (Darken & Goerger, 1999) and whether the cost is justified (though it is obvious that it is important for entertainment and seems to sell simulators). The computer game and movie/TV special effects industry has motivated not only new software techniques that take advantage of higher-speed computers and higher display resolutions but also new head-mounted displays, “data gloves,” and 3-D auditory display techniques that enable the immersion effect. The effect is enhanced not only by higher-resolution displays but also by the user’s being able to change the viewpoint, head orientation, and hand position relative to a touched object and to have the sensory pattern change according to expectation (Sheridan, 1992a).

In due course the technology will meet the criterion of the Turing test, whereby the observer will not be able to discriminate what is virtual from what is real. That poses serious concerns—for example, worries that children who spend hours playing violent and very realistic computer games will transfer that violent behavior into their real-world lives. Even at the seemingly more innocuous level of mechanized toys there is concern that the toys' use is too preprogrammed and children's development of imagination is inhibited (Sheridan, 1993).

The compelling nature of virtual reality is illustrated from a piece in *Wired Today.* After a recent three-day binge of playing the Japanese cult hit video game Katamari Damacy, Los Angeles artist Kozy Kitchens discovered that walking away from the game was not as easy as putting down her joystick. In the game, players push around what amounts to a giant tape ball, attempting to make the ball bigger by picking up any and all objects in its path. Kitchens found that her urge to keep picking things up was not so easy to shake.

"I was driving down Venice Boulevard," recalled her husband, Dan Kitchens, "and Kozy reached over and grabbed the steering wheel and for a moment was trying to yank it to the right. . . . (Then) she let go, but kept staring out her window, and then looked back at me kind of stunned and said, 'Sorry. I thought we could pick up that mailbox we just passed.'"

Though motorists and pedestrians shouldn't worry too much about rogue Katamari Damacy players, Kozy Kitchens' difficulty with separating her real-life consciousness from that of her game playing is all too common among hard-core gamers. It's so common, in fact, that game publishers might want to consider warning their customers that they may soon be unable to tell the difference between the game and reality. Frequent gamer Alfred Weisberg-Roberts said he often feels lingering effects after playing games like Animal Crossing, in which the point is to collect as many animals and bugs as possible from a wide variety of locations.

"Once, my girlfriend happened upon a tree . . . kind of like the round, thin trees in the game, and began to shake it—one in-game way of receiving money, goods, and bees," Weisberg-Roberts said. "When nothing fell from its branches, I think she quickly realized how this must have looked to the other hundred or so people in the park."

Mixed-Initiative Conflicts within Large-Scale Systems

Large-scale systems are characterized by many sensors, many different computers as well as people performing analyses and making decisions, and many actuators taking physical actions to implement those decisions. This usually means that there are multiple control loops, each trying to drive some variable to correspond with its reference input in spite of external disturbances. The problem is that within the physical process being controlled, these actions can be and often are coupled—meaning the action of one control loop appears as a disturbance to the other control loop. For example, in a robotic system one control loop may be programmed to drive the robot to move toward a target, and a second control loop is programmed to avoid obstacles. If the

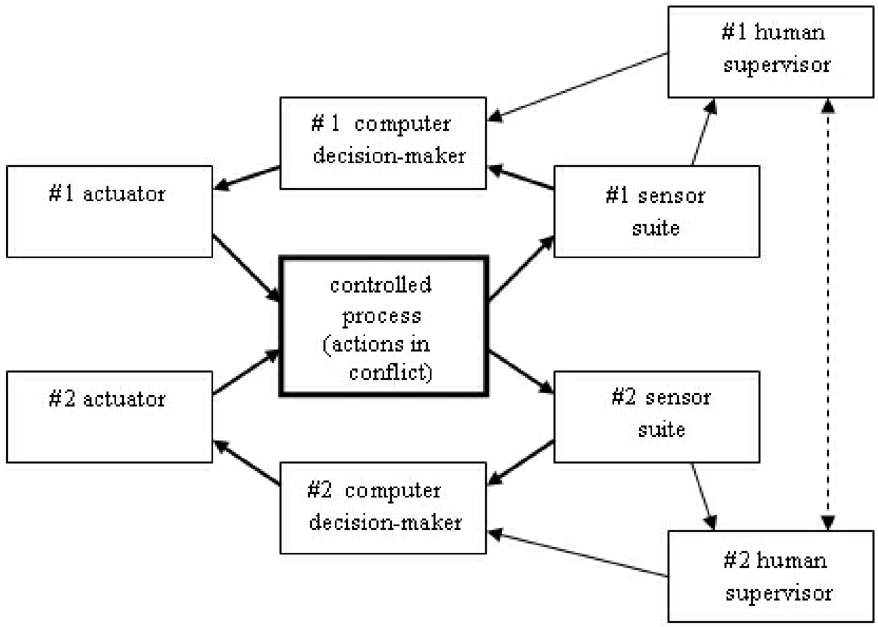

Figure 2.3. *Mixed-initiative conflicts within large-scale automation.*

robot runs into an obstacle while moving toward the target, there is an impasse. Figure 2.3 illustrates the idea; the heavy arrows designate the two automated control loops. This is commonly called a *mixed-initiative problem.*

Human supervisors may have access to the same sensed information as the automation they are overseeing but not to information from the other control loop. In other words, they see different sides of the elephant. If the two supervisors had access to each other's information, in theory they might be able to resolve the conflict, but the additional information just adds to their workload, and furthermore the communication between them (dashed line) is likely to be delayed, noisy, or nonexistent.

Another type of conflict arises when the actuators are automatically drawing from a common resource pool (or supplies, energy, budgeted money, etc.) and each supervisor is trying to fulfill his or her own goals independent of the others. This is what Hardin (1968) has called "the tragedy of the commons." Naturally, when there are more control loops and more people in the system, these problems are compounded.

Automation as a Cause of User Alienation

We use the term *alienation* to characterize a number of social and political issues concerning the individual. In an automated system the human is removed not only spatially but also temporally, functionally, and cognitively from the ongoing physical process he or she is directing. What the human does and thinks is likely to have different timing, different explicit form, and different logical content from what the automation does and thinks. Various authors have discussed the alienation problem (e.g., see Hofstede, 1994; Moray, 2000; Zuboff, 1988). Sheridan (1980) presented some components of alienation, as follows.

Threatened or actual unemployment. Organizations have become more efficient in their use of human supervisors: Fewer people can now supervise more machines. Furthermore, automation is becoming more able to detect its own failures, and in some cases it can even repair itself. The real threat has been unemployment of the unskilled and technologically illiterate. However, with the Internet breaking down geographic barriers, even the skilled and technologically literate are in danger of unemployment as high-skill jobs move to technologically sophisticated regions in Asia. The threat of technological displacement has moved up the chain, from the manual laborer and railroad brakeman to the computer system administrator, radiologist, and architect.

Erratic mental workload and dissatisfaction with work. Automation affects not only the nature but also the pace of work and may at times make that pace vary between extremes. This is the oft-cited "hours of boredom punctuated by moments of terror" syndrome.

Centralization of management control and loss of worker control. A result of automation and associated electronic technology is that management can secretly record and monitor workers. The mere possibility of being monitored in this way is often sufficient to produce worker anxiety, including fear that private data stored electronically may be accessed by persons other than those authorized. In some cases centralized monitoring may enhance productivity, but in other cases it may prove detrimental.

Desocialization. Interaction with computers is gradually replacing interaction with other people. As supervisory control systems are interconnected, the computer will mediate increasingly more of what interpersonal contact remains, as has already happened in many cases with e-mail, pagers, and associated software for management coordination and computer-supported cooperative work.

Deskilling. Skilled workers who are "promoted" to supervisory controllers (sometimes derogatorily referred to as "button pushers") may resent the transition. In part, this may be out of fear that when called on to take over and do the job manually, they may not be able to do so.

Intimidation of greater power. Automation encourages larger aggregations of interconnected equipment, higher speeds, greater complexity, and probably greater economic risk if something goes wrong and the supervisor doesn't take the appropriate corrective action. The human supervisor will be forced to assume increasingly more ultimate responsibility, although in most cases the responsibility probably should reside with some combination of the manager and the system designer.

Technological illiteracy. In the role of supervisory controller, the operator may lack technological understanding of how the computer and the rest of the complex technology do what they do. What is really going on with the communications and control software may be too specialized even for many technicians involved with the newer systems. The push toward increasing functionality and capability may mean that systems are becoming increasingly less visible and comprehensible to even the designers and maintainers.

Mystification and misplaced trust. Human operators of computer-based systems sometimes become mystified by and superstitious about the power of the computer, even seeing it as a kind of magic or a "big brother" authority figure. This leads naturally to misplaced trust.

Sense of not contributing. Though the efficiency and mechanical productivity of a new supervisory control system may far exceed that of an earlier manually controlled system, the operator may come to feel that with automation he or she is no longer the source of value added, no longer a significant contributor. The sense of personal productivity—what psychologist Erich Fromm (1995) called the *productive orientation*—is allegedly fundamental to humans' sense of self-worth. Without it, who are we?

Abandonment of responsibility. As a result of the factors just described, human supervisors of automation may eventually feel they are no longer responsible for what happens but that the computers are. A worker with his or her own set of hand tools or a simple, self-powered, but manually controlled machine—though he or she may sometimes place the blame for difficulties elsewhere—has a clear responsibility for use and maintenance of the tools or machine. When workers' actions are mediated by a powerful computer, however, the lines of responsibility are not so clear, and the workers may, in effect, abandon their responsibility for the task performed or the good produced, believing instead that it is in the "hands" of the computer.

Blissful enslavement. To many writers the worst form of alienation, the worst tragedy, occurs when a worker is happy to accept a role in which he or she is made to feel powerful but, in actuality, he or she is enslaved. Both Aldous Huxley's *Brave New World* and George Orwell's *1984* are famous for this theme of blissful enslavement.

Ease of Committing Violent Acts by Remote Control of Automation

Today's automation coupled with modern communication technology means that anyone can supervise an automatic machine from any arbitrary distance away—and without the source of control being apparent. This has been a boon to developers of unmanned space, aerial, undersea, and land vehicles, most of which have been developed for military use. In contrast to past warfare, in which the fighter risked his or her own life for the cause, today remote-control capabilities mean that violence can be committed anonymously by anyone without the perpetrator's even being aware of the result—for example, many innocents may be killed because a highly automated and even precisely targeted missile encountered circumstances on the ground that were unpredicted. No one likes terrorists with car bombs or bombs strapped to their bodies. But we can look forward to a day when violent acts or just insidious spying can be committed by unmanned vehicles and robots controlled by anyone able to acquire the communication and control technology, which will surely be getting smaller and cheaper. Social and moral responsibility increases as automation affects more people in more profound ways.

So What Should Not Be Automated?

At the end of this chapter we directly pose the question that is often asked of professionals concerned with humans and automation: What should not be automated, even though it is possible? Automation engineers—at least many of them—have the attitude that if it is technologically, economically, and legally feasible to automate, then do it; it is an exciting challenge (Bainbridge, 1983). Much of the foregoing discussion suggests reasons to go slow, to anticipate and evaluate the much more subtle problems that automation can bring. The human factors engineer is often regarded by the engineer as a worry-wart, a nay-sayer, a wet blanket in terms of risk taking and progress. But the assessment of human-related questions—potential effects on individual and social behavior, institutions, and culture—must be asked because, after all, the ultimate purpose of technology is to make life better for people (Hancock, 1996). Furthermore, these hard questions must be asked early—before the technological development is too far along, before the point of no return, or at least before the point where changes are very much more expensive than they would have been had they been made early.

There is a belief among many automation engineers that one can eliminate human error by eliminating the human operator. To the extent a system is made less vulnerable to operator error, it is made more vulnerable to designer error (Parasuraman & Riley, 1997). And given that the designer is also human, this simply displaces the locus of human error. In the end, automation is really human after all.

REFERENCES

Abbott, K., Slotte, S., Stimson, D., Amalberti, R. R., Bollin, G., Fabre, F., et al. (1996). *The interfaces between flightcrews and modern flight deck systems* (Report of the FAA Human Factors Team). Washington, DC: Federal Aviation Administration.

Aviation coding manual. (1998). Retrieved October 10, 2005, from http://www.ntsb.gov/aviation/codman_intro.htm

Aviation Safety Reporting System. (2005). Retrieved October 10, 2005, from http://asrs.arc.nasa.gov/

Bagheri, N., & Jamieson, G. A. (2004). Considering subjective trust and monitoring behavior in assessing automation-induced "complacency." In D. A. Vicenzi, M. Mouloua, & P. A. Hancock (Eds.), *Human performance, situation awareness, and automation: Current research and trends* (pp. 54–59). Mahwah, NJ: Erlbaum.

Bainbridge, L. (1983). Ironies of automation. *Automatica, 19,* 775–779.

Barnes, M., & Grossman, J. (1985). *The intelligent assistant concept for electronic warfare systems* (Technical Report NWC TP 5585). China Lake, CA: Naval Warfare Center.

Bennett, K., & Flach, J. (1992). Graphical displays: Implications for divided attention, focused attention, and problem solving. *Human Factors, 34,* 513–533.

Billings, C. E. (1997). *Aviation automation: The search for a human centered approach.* Mahwah, NJ: Erlbaum.

Billings, C. E., Lauber, J. K., Funkhouser, H., Lyman, G., & Huff, E. M. (1976). *NASA aviation safety reporting system* (Technical Report TM-X-3445). Moffett Field, CA: NASA Ames Research Center.

Billings, C. E., & Woods, D. D. (1994). Concerns about adaptive automation in aviation systems. In R. Parasuraman & M. Mouloua (Eds.), *Human performance in automated systems: Current research and trends.* (pp. 264–269). Mahwah, NJ: Erlbaum.

Byrne, E. A., & Parasuraman, R. (1996). Psychophysiology and adaptive automation. *Biological Psychology, 42,* 249–268.

Christoffersen, K., & Woods, D. D (2002). How to make automated system team players. In E. Salas (Ed.), *Advances in human performance and cognitive engineering research* (vol. 2, pp. 1–12). Amsterdam: Elsevier.

Comstock, J. R., & Arnegard, R. J. (1992). Multi-attribute task battery (NASA Technical Memorandum 104174). Hampton, VA: NASA Langley Research Center.

Craik, K. J. W. (1947). Theory of the human operator in control systems, I: The operator as an engineering system. *British Journal of Psychology, 38,* 56–61.

Darken, R. P., & Goerger, S. R. (1999). The transfer of strategies from virtual to real environments: An explanation for performance differences? In *Proceedings of Virtual Worlds and Simulation '99* (pp. 159–164). La Jolla, CA: Society for Computer Simulation International.

Degani, A. (2003). *Taming Hal: Designing interfaces beyond 2001.* New York: Palgrave MacMillan.

Dialogs on function allocation [Special issue]. (2000). *International Journal of Human-Computer Studies (Vol 52).*

Endsley, M., & Kaber, D. (1999). Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics, 42,* 462–492.

Erol, K., Hendler, J., & Nau, D. (1994). UMCP: A sound and complete procedure for hierarchical task network planning. In K. Hammond (Ed.), *AI planning systems: Proceedings of the 2nd International Conference* (pp. 249–254). Los Altos, CA: AAAI.

Fitts, P. M. (2005). Some basic questions in designing an air-navigation and air-traffic control system. In N. Moray (Ed.), *Ergonomics major writings* (Vol. 4, pp. 367–383). London: Taylor & Francis. (Washington DC, Reprinted from *Human engineering for an effective air navigation and traffic control system,* National Research Council, pp. 5–11, 1951, Washington, DC: National Academy Press.)

Fromm, E. (1995). *Escape from freedom.* New York: Holt.

Funk, K., Lyall, B., Wilson, J., Vint, R., Miemcyzyk, M., Suroteguh, C., et al. (1999). Flight deck automation issues. *International Journal of Aviation Psychology, 9,* 125–138.

Furukawa, H., Parasuraman, R., & Inagaki, T. (2003). Supporting system-centered view of operators through ecological interface design: Two experiments on human-centered automation. In *Proceedings of the Human Factors and Ergonomics Society* (pp. 567–571). Santa Monica, CA: Human Factors and Ergonomics Society.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton-Mifflin.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics: Speech acts* (Vol. 3, pp. 276–290). New York: Academic.

Hammer, J. M., & Small, R. L. (1995). An intelligent interface in an associate system. In W. B. Rouse (Ed.), *Human/technology interaction in complex systems* (Vol. 7, pp. 1–44). Greenwich, CT: JAI Press.

Hancock, P. A. (1996). Teleology for technology. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications* (pp. 461–497). Mahwah, NJ: Erlbaum.

Hancock, P. A., Chignell, M. H., & Lowenthal, A. (1985). An adaptive human-machine system. In *Proceedings of the IEEE Conference on Systems, Man and Cybernetics, 15* (pp. 627–629). Washington, DC: IEEE.

Hancock, P. A., & Scallen, S. F. (1996, October). The future of function allocation. *Ergonomics in Design* (October), 24–29.

Hardin, G. (1968). The tragedy of the commons. *Science, 162,* 1243–1248.

Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology, 4,* 1–26.

Hilburn, B., Jorna, P. G., Byrne, E. A., & Parasuraman, R. (1997). The effect of adaptive air traffic control (ATC) decision aiding on controller mental workload. In M. Mouloua & J. Koonce (Eds.), *Human-automation interaction* (pp. 84-91). Mahwah, NJ: Erlbaum.

Hofstede, G. (1994). *Cultures and organizations.* London: HarperCollins.

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

Inagaki, T. (2003). Adaptive automation: Sharing and trading of control. In E. Hollnagel (Ed.), *Handbook of cognitive task design* (pp. 221–245). Mahwah, NJ: Erlbaum.

James, H., Nichols, N., & Phillips, R. (1947). Manual tracking. In *Theory of servomechanisms.* New York: McGraw Hill.

Jamieson, G. A., & Vicente, K. J. (2005). Designing effective human-automation-plant interfaces: A control theoretic perspective. *Human Factors, 47,* 12–34.

Jordan, N. (1963). Allocation of functions between man and machines in automated systems. *Journal of Applied Psychology, 47,* 161–165.

Kaber, D. B., & Endsley, M. (2004). The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. *Theoretical Issues in Ergonomics Science, 5,* 113–153.

Kaber, D. B., & Riley, J. M. (1999). Adaptive automation of a dynamic control task based on workload assessment through a secondary monitoring task. In M. Scerbo & M. Mouloua (Eds.), *Automation technology and human performance: Current research and trends* (pp. 129–133). Mahwah, NJ: Erlbaum.

Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a team player in joint human-agent activity. *IEEE Intelligent Systems 19*(6), 91–95.

Kleinman, D. L., Baron, S., & Levison, W. H. (1970). An optimal control model of human response, Part 1. *Automatica, 63,* 357–369.

LaPorte, T. R. (1996). High reliability organizations: Unlikely, demanding, and at risk. *Journal of Contingencies and Crisis Management, 4*(2), 60–71.

Lee, J. D., & Moray, N. (1992). Trust, control strategies and allocation of function in human-machine systems. *Ergonomics, 35,* 243–1270.

Lee, J., & Moray, N. (1994). Trust, self confidence, and operators' adaptation to automation. *International Journal of Human-Computer Studies, 40,* 153–184.

Lee, J., & See, J. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors, 46,* 50–80.

Lee, J. D., Caven, D., Haake, S., & Brown, T. L. (2001). Speech-based interaction with in-vehicle computers: The effect of speech-based e-mail on drivers' attention to the roadway. *Human Factors, 43,* 631–640.

Lee, J. D., McGehee, D., Brown, T. L., & Reyes, M. (2002). Collision warning timing, driver distraction, and driver response to imminent rear end collision in a high-fidelity driving simulator. *Human Factors, 44,* 314–334.

Lewandowsky, S., Mundy, M., & Tan, G. P. (2000). The dynamics of trust: Comparing humans to automation. *Journal of Experimental Psychology: Applied, 6,* 104–123.

Llaneras, R. E. (2000). *NHTSA driver distraction internet forum.* Retrieved October 10, 2005, from http://www-nrd.nhtsa.dot.gov/departments/nrd-13/DriverDistraction.html

McRuer, D. T., & Jex, H. R. (1967). A review of quasi-linear pilot models. *IEEE Transactions on Human Factors in Electronics, HFE-4*(3), 231–249.

McRuer, D., & Krendel, E. (1959). The human operator as a servo system. *Journal of the Franklin Institute, 267,* 5–6.

Metzger, U., & Parasuraman, R. (2001). Automation-related "complacency": Theory, empirical data, and design implications. In *Proceedings of the Human Factors and Ergonomics Society 45th Annual Meeting* (pp. 463–467). Santa Monica, CA: Human Factors and Ergonomics Society.

Metzger, U., & Parasuraman, R. (2005). Automation in future air traffic management: Effects of decision aid reliability on controller performance and mental workload. *Human Factors, 47,* 35–49.

Meyer, J. (2001). Effects of warning validity and proximity on responses to warnings. *Human Factors, 43,* 563–572.

Miller, C. A. (2004). Human-computer etiquette [Special issue]. *Communications of the ACM, 37*(4).

Miller, C. A., Goldman, R., Funk, F., Wu, P., & Pate, B. (2004, June). A Playbook approach to variable autonomy control: Application for control of multiple, heterogeneous unmanned air vehicles. In *Proceedings of Forum 60, the Annual Meeting of the American Helicopter Society.* Alexandria, VA: AHS International.

Miller, C. A., & Hannen, M. D. (1999). The rotorcraft pilot's associate: Design and evaluation of an intelligent user interface for cockpit information management. *Knowledge-Based Systems, 12,* 443–456.

Miller, C. A., & Parasuraman, R. (in press). Designin for flexible interaction between humans and automation. *Human Factors*.

Miller, C. A., Pelican, M., & Goldman, R. (2000). "Tasking" interfaces for flexible interaction with automation: Keeping the operator in control. In *Proceedings of the Conference on Human Interaction with Complex Systems* (pp. 123–128). Urbana-Champaign, IL: HICS.

Miller, G. A. (1956). The magical number 7, plus or minus 2: Some limits on our capacity for processing information. *Psychological Review, 63,* 81–97.

Molloy, R., & Parasuraman, R. (1994). Automation-induced monitoring inefficiency: The role of display integration and redundant color coding. In M. Mouloua & R. Parasuraman (Eds.), *Human performance in automated systems: Current research and trends* (pp. 224–228). Mahwah, NJ: Erlbaum.

Moray, N. (1986). Monitoring behavior and supervisory control. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance* (vol. 2, pp. 40-1–40-51). New York: Wiley.

Moray, N. (1992). Flexible interfaces can promote operator error. In J. Kragt (Ed.), *Case studies in ergonomics* (pp. 49–64). London: Taylor & Francis.

Moray, N. (2000). Culture, politics and ergonomics. *Ergonomics, 43,* 858–868.

Moray, N. (2005). *Ergonomics major writings.* London: Taylor & Francis.

Moray, N., & Inagaki, T. (2001). Attention and complacency. *Theoretical Issues in Ergonomics Science, 1,* 354–365.

Moray, N., Inagaki, T., & Itoh, M. (2000). Situation adaptive automation, trust and self-confidence in fault management of time-critical tasks. *Journal of Experimental Psychology: Applied, 6*(1), 44–58.

Morrison, J. G., & Gluckman, J. P. (1994). Definitions and prospective guidelines for the application of adaptive automation. In M. Mouloua & R. Parasuraman (Eds.), *Human performance in automated systems: Current research and trends* (pp. 256–263). Mahwah, NJ: Erlbaum.

Mosier, K., Skitka, L. J., Heers, S., & Burdick, M. (1998). Automation bias: Decision making and performance in high-tech cockpits. *International Journal of Aviation Psychology, 8,* 47–63.

Muir, B. M. (1988). Trust between humans and machines, and the design of decision aids. In E. Hollnagel, G. Mancini, & D. D. Woods (Eds.), *Cognitive engineering in complex dynamic worlds* (pp. 71–84). London: Academic.

Nass, C., Moon, Y., Fogg, B. J., Reeves, B., & Dryer, D. C. (1995). Can computer personalities be human personalities? *International Journal of Human-Computer Studies, 43,* 223–239.

National Transportation Safety Board. (1973). *Eastern Air Lines, Inc., L-1011, N310EA, Miami, Florida, December 29, 1972* (AAR-73-14). Washington, DC: Author.

National Transportation Safety Board. (1998a). *Brief of accident NYC98FA020.* Washington, DC: Author.

National Transportation Safety Board. (1998b). *Safety recommendation letter A-98-3 through -5, January 21, 1998.* Washington, DC: Author.

Norman, D. A. (1990). The problem with automation: Inappropriate feedback and interaction, not "over-automation." *Philosophical Transactions of the Royal Society* (London), B237, 585–593.

Opperman, R. (1994). *Adaptive user support.* Mahwah, NJ: Erlbaum.

Parasuraman, R. (1993). Effects of adaptive function allocation on human performance. In D. J. Garland & J. A. Wise (Eds.), *Human factors and advanced aviation technologies* (pp. 147–157). Daytona Beach, FL: Embry-Riddle Aeronautical University Press.

Parasuraman, R. (2000). Designing automation for human use: Empirical studies and quantitative models. *Ergonomics, 43,* 931–951.

Parasuraman, R., Bahri, T., Deaton, J. E., Morrison, J. G., & Barnes, M. (1992). *Theory and design of adaptive automation in aviation systems* (Technical Report, Code 6021). Warminster, PA: Naval Air Development Center.

Parasuraman, R., & Byrne, E. A. (2003). Automation and human performance in aviation. In P. Tsang & M. Vidulich (Eds.), *Principles of aviation psychology* (pp. 311–356). Mahwah, NJ: Erlbaum.

Parasuraman, R., Galster, S., Squire, P., Furukawa, H., & Miller, C. A. (2005). A flexible delegation interface enhances system performance in human supervision of multiple autonomous robots: Empirical studies with RoboFlag. *IEEE Transactions on Systems, Man & Cybernetics 35,* 481–493.

Parasuraman, R., & Miller, C. (2004). Trust and etiquette in high-criticality automated systems. *Communications of the Association for Computing Machinery, 47*(4), 51–55.

Parasuraman, R., Molloy, R., & Singh, I. L. (1993). Performance consequences of automation-induced "complacency." *International Journal of Aviation Psychology, 3,* 1–23.

Parasuraman, R., & Mouloua, M. (1996). *Automation and human performance: Theory and applications.* Mahwah, NJ: Erlbaum.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse and abuse. *Human Factors, 39,* 230–253.

Parasuraman, R., Sheridan, T., & Wickens, C. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man and Cybernetics, SMC-30(3),* 286–297.

Prinzel, L. J., Freeman, F. G., Scerbo, M. W., Mikulka, P. J., & Pope, A. T. (2000). A closed-loop system for examining psychophysiological measures for adaptive automation. *International Journal of Aviation Psychology, 10,* 393–410.

Rasmussen, J. (1986). *Information processing and human-machine interaction.* Amsterdam: North-Holland.

Rasmussen, J., Pedersen, A.-M., & Goodstein, L. (1995). *Cognitive engineering: Concepts and applications.* New York: Wiley.

Reason, J. T. (1990). *Human error.* Cambridge, England: Cambridge University Press.

Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places.* New York: Cambridge University Press.

Riley, V. (1996). Operator reliance on automation: Theory and data. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications* (pp. 19–35). Mahwah, NJ: Erlbaum.

Rochlin, E., LaPorte, T., & Roberts, K. (1987, Autumn). The self-designing high reliability organization: Aircraft flight operation at sea. *Naval War College Review,* 76–91.

Rouse, W. B. (1976). Adaptive allocation of decision making responsibility between supervisor and computer. In T. B. Sheridan & G. Johannsen (Eds.), *Monitoring behavior and supervisory control* (pp. 295–306). New York: Plenum.

Rouse, W. B. (1980). *System engineering models of human-machine interaction.* Amsterdam: North-Holland.

Rouse, W. B. (1988). Adaptive aiding for human/computer control. *Human Factors, 30,* 431–438.

Sarter, N., Woods, D., & Billings, C. E. (1997). Automation surprises. In G. Salvendy (Ed.), *Handbook of human factors and ergonomics* (2nd ed., pp. 1926–1943). New York: Wiley.

Sarter, N., & Woods, D. D. (1995). How in the world did we ever get into that mode? Mode error and awareness in supervisory control. *Human Factors, 37,* 5–19.

Scerbo, M. W. (1996). Theoretical perspectives on adaptive automation. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications* (pp. 37–63). Mahwah, NJ: Erlbaum.

Scerbo, M. W. (2001). Adaptive automation. In W. Karwowski (Ed.), *International encyclopedia of ergonomics and human factors* (pp. 1077–1079). London: Taylor & Francis.

Scerbo, M., Freeman, F., Mikulka, P. J., Di Nocera, F., Parasuraman, R., & Prinzel, L. (2001). *The efficacy of physiological measures for implementing adaptive technology* (NASA Technical Memorandum). Hampton, VA: NASA Langley Research Center.

Senders, J. W. (1964). The human operator as a monitor and controller of multidegree of freedom systems. *IEEE Transactions on Human Factors in Electronics, HFE-5,* 1–6.

Shannon, C. E. (1947). Communication in the presence of noise. *Proceedings of the IRE, 37,* 10–22.

Sheridan, T. B. (1960). The human metacontroller. In *Proceedings of the Annual Conference on Manual Control.* Stamford, CT: Dunlap Associates.

Sheridan, T. B. (1976). Toward a general model of supervisory control. In T. B. Sheridan & G. Johannsen (Eds.), *Monitoring behavior and supervisory control* (pp. 271–282). Elmsford, NY: Plenum.

Sheridan, T. B. (1980, October). Computer control and human alienation. *Technology Review,* 60–73.

Sheridan, T. (1988). Trustworthiness of command and control systems. In *Proceedings of the International Federation of Automatic Control Symposium on Man-Machine Systems* (pp. 427–431). Elmsford, NY: Pergamon.

Sheridan, T. B. (1992a). Musings on telepresence and virtual presence. *Presence: Teleoperators and Virtual Environments, 1(1),* 120–126.

Sheridan, T. B. (1992b). *Telerobotics, automation, and human supervisory control.* Cambridge: MIT Press.

Sheridan, T. B. (1993). My anxieties about virtual environments. *Presence: Teleoperators and Virtual Environments, 2*(2), 141–142.

Sheridan, T. B. (1998). Allocating functions rationally between humans and machines. *Ergonomics in Design, 6*(3), 20–25.

Sheridan, T. B. (2000). Function allocation: Algorithm, alchemy, or apostasy? *International Journal of Human-Computer Studies, 52,* 203–216.

Sheridan, T. B. (2002). *Humans and automation: Systems design and research issues.* Santa Monica/New York: Human Factors and Ergonomics Society/Wiley.

Sheridan, T. B., & Ferrell, W. R. (1974). *Man-machine systems.* Cambridge: MIT Press.

Sheridan, T., & Parasuraman, R. (2000). Human vs. automation in responding to failures: An expected value analysis. *Human Factors, 42*, 403–407.

Sheridan, T. B., & Thompson, J. M. (1994). People vs. computers in medicine. In S. Bogner (Ed.), *Human error in medicine.* Mahwah, NJ: Erlbaum.

Sheridan, T. B., & Verplank, W. L. (1978). *Human and computer control of undersea teleoperators* (Man-Machine Systems Lab Report). Cambridge: Massachusetts Institute of Technology.

Sherry, L., & Polson, P. G. (1999). Shared models of flight management system vertical guidance. *International Journal of Aviation Psychology, 9,* 139–153.

St. John, M., Kobus, D. A., Morrison, J. G., & Schmorrow, D. (2004). Overview of the DARPA augmented cognition technical integration experiment. *International Journal of Human-Computer Interaction, 17,* 131–149.

Swets, J. (1996). *Signal detection theory and ROC analysis in psychology and diagnostics.* Mahwah, NJ: Erlbaum.

Vicente, K. J. (2002). Ecological interface design: Progress and challenges. *Human Factors, 44,* 62–78.

Vicente, K., & Rasmussen, J. (1992). Ecological interface design: Theoretical foundations. *IEEE Transactions on Systems, Man and Cybernetics, SMC-22,* 589–606.

Wei, Z., Macwan, A., & Wierenga, P. (1998). A quantitative measure for degree of automation, and its relation to system performance. *Human Factors, 40,* 277–295.

Weiss, G. (Ed.). (1999). *Multi-agent systems.* Cambridge: MIT Press.

Wickens, C. D., & Hollands, J. G. (2000). *Engineering psychology and human performance.* Upper Saddle River, NJ: Prentice Hall.

Wickens, C. D., Mavor, A., Parasuraman, R., & McGee, J. (1998). *The future of air traffic control: Human operators and automation.* Washington, DC: National Academy Press.

Wiegmann, D. A., & Shappell, S. A. (1997). Human factors analysis of post-accident data: Applying theoretical taxonomies of human error. *International Journal of Aviation Psychology, 7,* 67–81.

Wiener, E. L. (1982). *Human factors of advanced technology: "Glass cockpit" transport aircraft* (NASA Contractor Report 177528). Moffett Field, CA: NASA Ames Research Center.

Wiener, E. L. (1988). Cockpit automation. In E. Wiener (Ed.), *Human factors in aviation.* San Diego, CA: Academic.

Wiener, E. L., & Curry, R. E. (1980). Flight-deck automation: Promises and problems. *Ergonomics, 23,* 995–1011.

Wiener, N. (1964). *God and Golem, Inc.* Cambridge: MIT Press.

Wilson, G. F., & Russell, C. A. (2003). Real-time assessment of mental workload using psychophysiological measures and artificial neural networks. *Human Factors, 45,* 635–643.

Woods, D. D. (1996). Decomposing automation: Apparent simplicity, real complexity. In R. Parasuraman & M. Mouloua (Eds.), *Automation technology and human performance.* Mahwah, NJ: Erlbaum.

Zuboff, S. (1988). *In the age of the smart machine: The future of work and power.* New York: Basic Books.