

Implications of Adaptive vs. Adaptable UIs on Decision Making: Why “Automated Adaptiveness” is Not Always the Right Answer

Christopher A. Miller, Harry B. Funk, Robert P. Goldman, John Meisner, Peggy Wu

Smart Information Flow Technologies
Minneapolis, MN
{cmiller, hfunk, rgoldman, jmeisner, pwu@sift.info}

Abstract

In this paper, we begin by contrasting the “adaptive” automation systems prevalent in Augmented Cognition research and development with “adaptable” systems that provide the same range of modifications and behaviors but place control over those adaptations in the hands of the human operator. While there are many obvious reasons to seek adaptive systems, there are also some less well-known consequences to high levels of automation control in human-machine systems. We review the literature on human interaction with high levels of automation in complex and critical work environments to illustrate various problematic effects including loss of situation awareness, poorly tuned trust, skill degradation, unbalanced mental workload, lack of user acceptance and, most importantly, poorer overall human + machine system performance and decision making. We illustrate how these effects may be mitigated by maintaining a more adaptable approach to automation use and suggest an approach to flexible adaptability, our *Playbook*TM delegation approach, which seeks to provide the best of both approaches.

1 Introduction

Opperman (1994) distinguishes between “adaptive” and “adaptable” systems. In either case, flexibility exists within the system to adapt to changing circumstances, but his distinction centers on who is in charge of that flexibility. For Opperman, an adaptable system is one in which the flexible control of information or system performance automation is resides in the hands of the user; s/he must explicitly command, generally at run time, the changes which ensue. In an adaptive system, by contrast, the flexibility in information or automation behavior is controlled by the system. It is as if Opperman is implying (though not explicitly defining) a kind of “meta-automation” which is present and in control of the degrees of freedom and flexibility in information and performance automation subsystems in an adaptive system, but which is absent (and is replaced by human activities) in an adaptable one. It is unclear whether the Augmented Cognition community consistently uses Opperman’s terms or makes his distinction, but it would seem that, in the majority of cases at least, when the phrases “adaptive system”, “adaptive user interface” and “adaptive automation” are used in this community, they are used in Opperman’s sense of a machine system which controls flexibility in information and performance subsystems, albeit in the service of the human.

Adaptive systems tend to have some distinct advantages over adaptable ones in terms of their impact on human + machine system decision making, and these advantages make them useful in a wide range of military and commercial contexts. By effectively delegating the “meta-automation” control tasks to another agent (that is, off loading them from the human), adaptive systems can frequently achieve greater speed of performance, reduced human workload, more consistency, a greater range of flexibility in behaviors and can require less training time than do human-mediated adaptable systems. On the other hand, by taking the human operator out of that portion of the control “loop”, adaptive systems run some risks with regards to decision making that adaptable ones generally do not. Since this community is, perhaps, more familiar with the advantages of adaptive systems than the risks and the complimentary advantages of adaptable approaches, we will concentrate on the risks and disadvantages of adaptive systems below.

2 Disadvantages of Fully Adaptive Systems

Even when automation is fully competent to perform a function without human intervention or monitoring, there may still be reasons to retain human involvement. Increasing the level of automation of a given task and/or giving more tasks to automation, necessarily means decreasing the human role and involvement in those tasks. A wealth of research over the past 20 years points to some distinct disadvantages stemming from reduced human engagement and, by contrast, of advantages to be obtained from maintaining higher levels of human involvement in tasks—a characteristic of adaptable systems. A growing body of research has examined the characteristics of human operator interaction with automation and described the human performance costs that can occur with certain forms of automation (Amalberti, 1999; Bainbridge, 1983; Billings, 1997; Parasuraman & Riley, 1997; Parasuraman, Sheridan, & Wickens, 2000; Rasmussen, 1986; Sarter, Woods, & Billings, 1997; Satchell, 1998; Sheridan, 1992; Wickens, Mavor, Parasuraman, & McGee, 1998; Wiener & Curry, 1980). These performance problems are briefly summarized here.

2.1 Reduced Situation and System Awareness

High levels of automation, particularly of decision-making functions, may reduce the operator's awareness of certain system and environmental dynamics (Endsley & Kiris, 1995; Kaber, Omal, & Endsley, 1999). Humans tend to be less aware of changes in environmental or system states when those changes are under the control of another agent (whether that agent is automation or another human) than when they make the changes themselves (Wickens, 1994). Endsley and Kiris (1995) used an automobile driving decision making task with 5 levels of automation ranging from fully manual to fully autonomous and then asked subjects a series of situation awareness questions. In spite of the fact that there were no distracter tasks and subjects had no responsibilities other than either making driving decisions or monitoring automation in the making of them, results showed situation awareness for the rationale behind decisions was highest in the fully manual condition, intermediate for the intermediate automation levels and lowest for the full automation condition. Studies by Endsley and colleagues, suggest that a moderate level of decision automation providing decision support but leaving the human remains in charge of the final choice of a decision option is optimal for maintaining operator situation awareness.

Mode errors are another example of the impact of automation on the user's awareness of system characteristics (Sarter & Woods, 1994). A mode refers to the setting of a system in which inputs to the system result in outputs specific to that mode but not to other modes. Mode errors can be relatively benign when the number of modes is small and transitions between modes do not occur without operator intervention. For example, in using a remote controller for a TV/VCR, it is commonplace to make mistakes like pressing functions intended to change the TV display while the system is in the VCR mode, or vice versa. When the number of modes is large, however, as in the case of an aircraft flight management system (FMS), the consequences of error can be more significant. Mode errors arise when the operator executes a function that is appropriate for one mode of the automated system but not the mode that the system is currently in (Sarter et al., 1997). Furthermore, in some systems mode transitions can occur autonomously without being immediately commanded by the operator, who may therefore be unaware of the change in mode. If the pilot then makes an input to the FMS which is inappropriate for the current mode, an error can result. Several aviation incidents and accidents have involved this type of error (Billings, 1997; Parasuraman & Byrne, 2002).

2.2 Trust, Complacency, and Over-Reliance

Trust is an important aspect of human interaction with automation (Lee & Moray, 1992, 1994). Operators may not use a well-designed, reliable automated system if they believe it to be untrustworthy. Conversely, they may continue to rely on automation even when it malfunctions and may not monitor it effectively. Both phenomena have been observed (Parasuraman & Riley, 1997). Mistrust of automation, especially automated alerting systems, is widespread in many work settings because of the problem of excessive false or nuisance alarms.

The converse problem of excessive trust or complacency has also been documented. Several studies have shown that humans are not very good at monitoring automation states for occasional malfunctions if their attention is occupied with other manual tasks (Parasuraman, 1993). Parasuraman, Molloy and Singh (1993) showed evidence of increased complacency among users of highly, but not completely, reliable automation in laboratory settings. Metzger and Parasuraman (2001) reported similar findings for experienced air traffic controllers using decision aiding automation. In these studies, users effectively grant a higher level of automation to a system than it was designed to

support by virtue of coming to accept automatically the system's recommendations or processed information even though the system sometimes fails. Riley (1994) documents a similar phenomenon, overreliance on automation, by trained pilots. In an experiment where the automation could perform one of a pair of tasks for the operator, but would occasionally fail, almost all of a group of students detected the failure and turned the automation off, while nearly 50% of the pilots failed to do so. While it is impossible to conclude that pilots' increased experience with (albeit, reliable) automation is the cause for this overreliance, it is tempting to do so.

Over-reliance on automation can also be manifest as a bias in reaching decisions. Human decision-makers exhibit a variety of biases in reaching decisions under uncertainty. Many of these biases reflect decision heuristics that people use routinely as a strategy to reduce the cognitive effort involved in solving a problem (Wickens, 1992). Heuristics are generally helpful but their use can cause errors when a particular event or symptom is highly representative of a particular condition and yet is extremely unlikely. Systems that automate decision-making may reinforce the human tendency to use heuristics and result in a susceptibility to "automation bias" (Mosier & Skitka, 1996). Although reliance on automation as a heuristic may be an effective strategy in many cases, over-reliance can lead to errors, as in the case of any decision heuristic. Automation bias may result in omission errors, when the operator fails to notice a problem or take an action because the automation fails to inform the operator to that effect. Commission errors occur when operators follow an automated directive that is inappropriate. Evidence for both types of errors was reported.

2.3 Skill Degradation

If placing automation in a higher, more encompassing role can result in complacency and loss of situation awareness, it is perhaps not surprising that it can also result in skill degradation if allowed to persist over time. The pilots of increasingly automated aircraft feared this effect with regards to psychomotor skills such as aircraft attitude control (Billings, 1997), but it has also been demonstrated to occur for decision making skills (Kaber, et al., 1999). In both cases, the use of an intermediate, lower level of automated assistance proved to alleviate skill degradation assuming the skills had been learned in the first place.

2.4 Unbalanced Mental Workload

Automation can sometimes produce extremes of workload, either too low or too high. That automation can increase workload one of the "ironies of automation", because many automated systems when first introduced are touted as workload-saving moves, and the technical justification for automation often is that it reduces mental workload and hence human error. But this does not always occur. First, if automation is implemented in a "clumsy" manner, e.g., if executing an automated function requires extensive data entry or "reprogramming" by human operators at times when they are very busy, workload reduction may not occur (Wiener, 1988). Second, if engagement of automation requires considerable "cognitive overhead," (Kirlik, 1993), i.e. extensive cognitive evaluation of the benefit of automation versus the cost of performing the task manually, then users may experience greater workload in using the automation. Alternatively, they may decide not to engage automation. (This is, of course, an even greater risk for adaptable automation than for adaptive automation.) Finally, if automation involves a safety-critical task, then pilots may continue to monitor the automation because of its potential unreliability. As Warm, Dember, and Hancock (1996) have shown, enforced monitoring can increase mental workload, even for very simple tasks. Thus any workload benefit due to the allocation of the task to the automation may be offset by the need to monitor the automation.

2.5 Performance Degradation

Most significantly, intermediate levels of human involvement in tasks can produce better overall performance of the human + machine system than either full manual or full automation levels, especially when human and automation roles are well structured and complimentary. In an experiment involving commercial aircraft navigation and route planning, Layton, Smith and McCoy (1994) provided human operators with one of three levels of automation support. In a 'sketching only' condition, (a highly manual approach), operators were required to create route plans using a map-based interface entirely on their own. The system would provide feedback about the feasibility of the human-proposed route in terms such as fuel loading, time of arrival and recommended altitudes. At the other end of the spectrum, a very high level of automation was provided by a system that automatically recommended a 'best' route according to its optimization criteria to the pilot. This 'full automation' mode was capable of providing supporting information about its recommended route plan, and of evaluating suggested alternatives, but only in response to explicit requests from the user. An intermediate level placed the automation in the role of supporting the

user-initiated route planning at a higher level of functionality. The user could ask for a route with specific characteristics (e.g., by way of Denver, avoiding turbulence greater than class 2, etc.) and have the system provide its best route that met such constraints. Each level was cumulative in that the user in the 'full automation' mode could choose to interact in full automation, intermediate or sketching only modes, or could switch between them.

Using this paradigm, Layton et al. found that humans in the intermediate and high automation conditions frequently explored more potential routing alternatives than they did in the highly manual condition; especially when the problem was complex and the range of potential considerations were large. In the sketching only (highly manual) condition, the process of arriving at a route was too difficult for the user to be able to try many alternatives consistently and fully. On the other hand, in the highly automated condition, users tended to accept the first route suggested by the automation without exploring it or its alternatives deeply. Even when they did explore it, the system's recommendation tended to narrow and bias their search. This outcome is similar to that seen in previous studies of complacency and automation bias. Users also tended to check the route provided for obvious mistakes rather than do the work of generating a route on their own to see if the computer's route was similar to the one they would have preferred. Users tended to take more factors into account more fully in their reasoning and the routing options selected were better in the intermediate automation condition. Particularly in trials when the automated route planning capabilities were suboptimal (e.g., because they failed to adequately consider uncertainty in future weather predictions), the intermediate level of automation produced better overall solutions. Layton, et al. suggest this was because users were better aware of the situation and hence, better able to both detect problems in the automation's recommended path, and to explore a range of alternatives quickly.

2.6 Decreased User Acceptance

Finally, one additional reason for preferring automation at the intermediate levels may be operator preference. Our own research and experience has shown that as automation begins to encroach on previously human-held tasks it suffers from a basic sociological problem: human operators want to remain in charge. This is probably particularly true of highly trained and skilled operators of complex, high-criticality systems such as aircraft, military systems, process control and power generation. For example, in developing the Rotorcraft Pilot's Associate (Miller, 1999), we interviewed multiple pilots and designers to develop a consensus list of prioritized goals for a "good" cockpit configuration manager. In spite of offering an advanced, automated aid capable of inferring pilot intent and managing information displays and many cockpit functions to conform to that intent, two of the top three items on the consensus list were "Pilot remains in charge of task allocation" and "Pilot remains in charge of information presented."

Similarly, Vicente (1999) sites examples of human interactions with even such comparatively mundane and low-risk automation as deep fryer timers in fast food restaurants (Bisantz, Cohen, Gravelle, & Wilson, 1996)., illustrating how human operators can become frustrated when they are forced to interact with automation which removes their authority to do their jobs in the best way they see fit. This review goes on to summarize extensive findings by Karasek and Theorell (1990) showing that jobs in which human operators have high psychological demands coupled with low decision latitude (the ability to improvise and exploit one's skills in the performance of a job) lead to higher incidences of heart disease, depression, pill consumption, and exhaustion.

3 A Tradeoff Space for Automation Effects

The above results, taken together, imply that applying sophisticated, adaptive and intelligent automation to manage information flow and equipment behavior for human consumers in complex systems and domains is not a panacea. Users in complex, high consequence domains are very demanding and critical of automation that does not behave according to their standards and expectations, and it has proven difficult to create systems that are correct enough to achieve user acceptance. Worse, as implied above, overly automated systems may well reduce the overall performance of the human + machine system if they do not perform perfectly both because they can reduce the human operator's awareness of the situation, degrade his/her skills and minimize the degree to which s/he has thought about alternate courses of action when the automation fails.

The tradeoff is not a simple two-way relationship between human workload and the degree to which human tasks are given to the automation, as is suggested above. Instead, we have posited a three-way relationship between three factors as illustrated in Figure 1:

1. the *workload* the human operator experiences in interacting with the system to perform tasks—workload that can be devoted to actually (“manually”) executing tasks or to monitoring and supervising tasks, or to selecting from various task performance options and issuing instructions to subordinates in various combinations,
2. the overall *competency* of the human + machine system to behave in the right way across a range of circumstances. Competency can also be thought of as “adaptiveness”—which is not to say the absolute range of behaviors that can be achieved by the human + machine system, but rather the range of circumstances which can be correctly “adapted to”, in which correct behavior can be achieved.
3. the *unpredictability* of the machine system to the human operator—or the tendency for the system to do things in ways other than expected/desired by the human user (regardless of whether those ways were

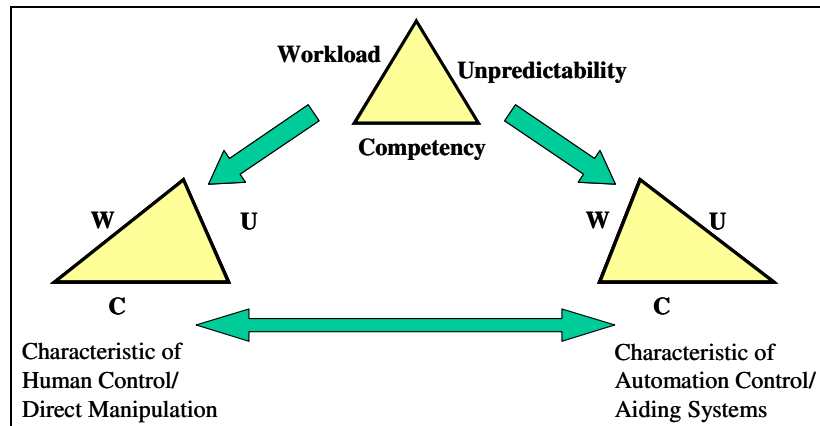


Figure 1. Three-way trade off space and its implications for a spectrum of control alternatives.

technically right). Unpredictability can be mitigated through good design, through training and through better (more transparent) user interfaces, but some degree of unpredictability is inherent whenever tasks are delegated and is a necessary consequence of achieving reductions in workload from task delegation.

As can be seen in Figure 1, human controlled, “direct manipulation” systems (Schneiderman, 1997) are those in which an increased share of the responsibility for achieving a given level of competency has been given to the human. Such systems necessarily entail more workload for the human, but mitigate uncertainty to the human by keeping him or her “in charge” of how things are done. By contrast, automation-controlled systems are those in which the responsibility for achieving the same level of competency has been placed in the hands of automation. Such systems characteristically reduce workload for the human operator, but only at the cost of greater unpredictability. It is possible to achieve a given level of competency through either an expansion in workload or an expansion in unpredictability—or various mixes in between. One implication of this three way relationship is that it is generally impossible to achieve both workload reduction and perfect predictability in any system which must adapt to complex contexts. Another important implication is that the endpoints of full automation vs. full human control describe a spectrum with many possible alternatives in between. Each alternative represents a different mix of task execution and monitoring activities that are the responsibility of the human operator versus that of the system automation. The spectrum of alternatives that result is roughly equivalent to the spectrum of choices that lies




Figure 2. Augmented Cognition technologies (and their like) enable real-time, dynamic flexibility in adapting a design across the spectrum.

between Direct Manipulation (Schneiderman, 1997) and Intelligent Agent interfaces (Maes, 1994).

At issue, however, is who decides what that mix of automation and human responsibility should be, and when. Traditionally, designers have made that decision statically at design time and called it “function allocation”. Increasingly, however, sophisticated world- and operator-sensing technologies (such as Augmented Cognition technologies) are making it possible to adapt that allocation during execution time in a highly dynamic fashion. The effect is as illustrated in Figure 2—the ability to dynamically choose where in the spectrum of human vs. automation task allocation the system should be “designed” at any point in time. The question that this capability raises, however, is where the decision about how to allocate tasks or functions should reside: with the human operator (making the resulting system an “adaptable” one in Opermann’s sense) or with automation (making it an “adaptive” one).

In general, the goals of making that decision an “adaptive” one are well understood—decreased workload for human operators, improved human + machine system performance through making better use of both more reliable automation capabilities and more efficient use of human resources. Nevertheless, as we saw above, placing the control of the decision about how to adapt in the hands of automation can have unexpected and undesirable consequences on both human acceptance and overall human + machine system behavior. While retaining that decision making function in the human’s hands has drawbacks of its own, we tend to favor more of an “adaptable” approach to interacting with dynamic automation, though the timing and even the scope of that decision can also be adapted as needed by the human operator. Below, we will describe the advantages we anticipate for more adaptable systems, while concluding with a brief description of an interaction architecture we are developing to allow the human operator to control the degree of authority s/he delegates to automation very dynamically and flexibly throughout the context of use.

4 Proposed Effects of Increasing Adaptability

There is reason to believe that adaptable systems do not suffer from the same set of problems as those described above. By keeping the operator in active charge of how much and what kind of automation to use when, we keep him or her more actively “in the loop” and, hence, more aware of how the system is or should be performing. By maintaining that level of involvement over automation behaviors, we would expect the following effects, each a mitigation or amelioration of the detrimental effects of high levels of automation as described above:

- Increased awareness of the situation and of system performance
- Requiring operators to make decisions about when to use automation, and to instruct automation in what behaviors to exhibit should produce better tuning of trust and better automation reliance decisions
- Allowing the user to perform tasks when needed or desired should keep skills more active and decrease the effects of skill degradation
- Allowing users more control over how much automation to use when will put them in a better position to manage their mental workload and keep it balanced
- If users can make good judgments (or simply better judgments than adaptive automation) about how much automation to use when to best compliment their workload, skills and capabilities, then a more nearly optimized mix of human and automation performance should result, along with the avoidance of performance degradation effects associated with full automation.
- Leaving the user in charge of when and how to use automation will enhance his/her sense of remaining in charge of automation performance, leading not only to a greater degree of acceptance, but also to a sense of being *primarily* responsible for overall task performance—in turn leading to greater attention and concern for the situation and all aspects of system performance.

Of course, adaptable systems have their own set of strengths and weaknesses. While an adaptable system would be expected to provide the benefits described above, it would suffer from increased workload on the part of the human operator and, perhaps, reduced overall task capacity for the human + machine system due to that workload. Indeed, while the human would be expected to have greater awareness of those aspects of the system to which s/he attended in an adaptable system, it might well be the case that fewer aspects/tasks could be attended to overall due to the attentional requirements placed on each task.

In fact, the work of Kirlik (1993) illustrates some of the downsides to adaptable automation. Kirlik developed a UAV simulation in which a human operator was in charge of manually controlling UAVs to have them visit various scoring points while simultaneously flying and navigating his/her own simulated helicopter. The pilot's own helicopter could be flown either manually or by putting it in autopilot mode. In this sense, the piloting of the own ship was an "adaptable" automation task in Opperman's terms. Kirlik performed a Markov Decision Process analysis to determine where decisions to use the autopilot would be optimal given a range of assumptions about (1) how much better or worse the human pilot was at the task than the autopilot, (2) how much time it took to engage and disengage the autopilot, and (3) the degree of inefficiency (as represented by a penalty for non-performance) for not having tasked the UAVs. The results of Kirlik's mathematical analysis showed distinct regions in which deciding to use the auto-pilot would and would not be optimal and, especially, showed that the effects of decision and execution time could eat into the effective performance of automation—implying both that the task of managing automation adds to the human's workload and may make it "more trouble than it's worth" to activate automation, and further implying that if that management task were to be successfully automated (as is the goal for adaptive automation) then there would be a greater likelihood of obtaining value from other forms of task automation. More importantly, in an experiment involving graduate students in the role of pilot, Kirlik found that subjects regularly avoided the least optimal strategies, but were inconsistent in their ability to find the most optimal ones.

In short, adaptable automation is no more a panacea than is adaptive automation. Humans may not always have the time or the expertise to choose the best forms of automation or the best times to use automation. Fortunately, however, we are not required to make hard choice between the adaptive and adaptable alternatives. We must strive to design a proper relationship between human operators and their automation that allows both parties to share responsibility, authority and autonomy over many work behaviors in a safe, efficient and reliable fashion. We are working on developing an approach that, while fundamentally adaptable in nature, nevertheless allows the human to "delegate" large spheres of authority within which adaptive behaviors are welcome and expected. This approach will be outlined in the next section below.

5 An Approach to Integrating Adaptive and Adaptable Systems

We have been exploring an approach to human interaction with complex automation that we call "delegation" because it is patterned on the kinds of interactions that a supervisor can have with an intelligent, trained subordinate. Human task delegation within a team or organizational setting is, overall, an adaptable system, in Opperman's sense, since the human supervisor can choose which tasks to hand to a subordinate, can choose what and how much to tell the subordinate about how (or how not) to perform the subtasks s/he is assigned, can choose how much or how little attention to devote to monitoring, approving, reviewing and correcting task performance, etc. Nevertheless, there can be substantial elements of adaptive system behavior within a delegation interaction in two ways. First, within tasks delegated to the subordinate (in this case, to automation), the subordinate may have full or partial authority to determine how to perform those tasks—including, potentially, what information to report to and even what tasks to ask from the supervisor. Second, there are opportunities for a good subordinate to take initiative within delegated roles to suggest tasks that need to be done and information that needs to be attended to—perhaps even supplying recommendations for who should do those tasks.

In work developing an interface for Uninhabited Combat Air Vehicles (UCAVs), we have explored a method of interacting with automation that attempts to more closely emulate human delegation relationships. In brief, our solution is to allow human operators to interact with advanced automation flexibly at a variety of automation levels and on a task-by-task basis. This allows the operator to smoothly adjust the 'amount' of automation s/he uses and the level at which s/he interacts with the hierarchy of tasks or functions to be performed depending on such variables as time available, workload, criticality of the decision, degree of trust, etc—variables known to influence human willingness and accuracy in automation use (Parasuraman and Riley, 1997).

This does not eliminate the dilemma presented in Figure 1, but it mitigates it by allowing operators to choose various points on the spectrum (as illustrated in Figure 2) for interaction with automation. The fundamental tradeoff between workload and unpredictability remains, but the operator is now put in charge of choosing a point in that tradeoff space. This strategy follows both Rasmussen's (Rasmussen & Goodstein, 1987) and Vicente's (1999) advice that operators should be allowed to 'finish the design' of the system at the time, and in the context, of use. This approach allows the user more control and authority over how and when s/he interacts with automation—and how that automation behaves. Therefore, it should address the desire to remain in charge that operators feel.

The trick, of course, is to design such systems so that they avoid two problems. First, they must make achievable the task of commanding automation to behave as desired without excessive workload. Second, they must ensure that resulting commanded behavior does, in fact, ensure safe and effective overall human + machine system behavior. We have created a design metaphor and system architecture that addresses these two concerns. Our approach to enabling, facilitating and ensuring correctness from a delegation interface, we call a Playbook™—because it is based on the metaphor of a sports team’s book of approved plays, with appropriate labels for efficient communication and a capability to modify, constrain, delegate and invent new plays as needed and as time permits.

We have written extensively about the Playbook architecture and applications elsewhere (Miller, Pelican and Goldman, 2000; Miller, et al., 2004; Miller and Parasuraman, submitted) and will not repeat that work here. Instead, we will conclude by pointing out three important attributes of delegation systems relevant to the creation of Augmented Cognition systems and the integration of the two:

1. The supervisor in a human-human delegation setting, and the operator of our Playbook™, maintains the overall position of authority in the system. It is not just that subordinates react to their perceptions of his/her intent, but rather they take explicit instructions from him/her. Subordinates may be delegated broad authority to make decisions about how to achieve goals or perform tasks, but this authority must come from the supervisor and not be taken autonomously by the subordinate. The act of delegation/instruction/authorization is, itself, important because it is what keeps the human supervisor “in the loop” about the subordinates activity and authority level. While it costs workload, if the system is well-tuned and the subordinate is competent, then that workload is well spent and results in net payoffs.
2. Even within its delegated sphere of authority, the subordinate does well to keep the supervisor informed about task performance, resource usage and general situation assessment. Beyond simple informing, the delegation system should allow some interaction over these parameters—allowing the supervisor to intervene to correct the subordinate’s assumptions or plans on these fronts. The supervisor may choose not to do so, but that should be his/her choice (and again, the making of that choice will serve to enhance awareness, involvement, empowerment and, ideally, performance).
3. There remain substantial opportunities for Augmented Cognition technologies to improve the competency of subordinate systems, and the ability for the human supervisor to maintain awareness of and control over them. Chief among these are methods to improve the communication of plans and recommendations between human and machine systems, to improve negotiation of plans (and plan revisions) so as to take best advantage of the strengths and weaknesses of both human and machine participants, and to provide plans, recommendations and other information when it is needed so as to improve uptake. Note that this last opportunity must be subservient to the first described above—the human supervisor should remain in charge of information presentation. A good subordinate must know when information or plans beyond what the supervisor has requested will be useful and valuable—but s/he must also know when and how to present them so as not to interrupt the supervisor’s important ongoing thoughts and actions.

In short, after years of attempting to design purely adaptive systems, in Opperman’s sense, we are skeptical about their utility in high complexity and high criticality domains. Instead, we opt for a more nearly adaptable approach that leaves the decision about when and what kind of automation to be used in the hands of a human operator/supervisor. Nevertheless, Augmented Cognition technologies have an important role to play in both types of systems.

6 Acknowledgements

The authors would like to thank Raja Parasuraman for his help in compiling earlier versions of this paper.

7 References

- Amalberti, R. (1999). Automation in aviation: A human factors perspective. In D. Garland, J. Wise & V. Hopkin (Eds.), *Handbook of aviation human factors* (pp. 173-192). Mahwah, NJ: Erlbaum.
- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19, 775-779.

- Billings, C. (1997). *Aviation automation: The search for a human-centered approach*. Mahwah, NJ: Erlbaum.
- Bisantz, A., Cohen, S., Gravelle, M., & Wilson, K. (1996). "To cook or not to cook: A case study of decision aiding in quick-service restaurant environments," Georgia Institute of Technology, College of Computing, Atlanta, GA, Report No. GIT-CS-96/03.
- Endsley, M. & Kiris, E. (1995). The out-of-the-loop performance problem and level of control in automation. *Human Factors*, 37, 381-394.
- Kaber, D., Omal, E. & Endsley, M. (1999). Level of automation effects on telerobot performance and human operator situation awareness and subjective workload. In M. Mouloua & R. Parasuraman, (Eds.), *Automation technology and human performance: Current research and trends*, (pp. 165-170). Mahwah, NJ: Erlbaum.
- Karasek, R. & Theorell, T. (1990). *Healthy work: Stress, productivity, and the reconstruction of working life*. New York: Basic Books.
- Kirlik, A. (1993). Modeling strategic behavior in human-automation interaction: Why an 'aid' can (and should) go unused. *Human Factors*, 35, 221-242.
- Layton, C., Smith, P., & McCoy, E. (1994). Design of a cooperative problem solving system for enroute flight planning: An empirical evaluation. *Human Factors*, 36(1), 94-119.
- Lee, J., & Moray, N. (1992). Trust, control strategies, and allocation of function in human-machine systems. *Ergonomics*, 35, 1243-1270.
- Lee, J., & Moray, N. (1994). Trust, self-confidence, and operators' adaptation to automation. *International Journal of Human-Computer Studies*, 40, 153-184.
- Maes, P. (1994). Agents that reduce work and information overload. *Communications of the Association for Computing Machinery*, 37(7), 31-40.
- Metzger, U. & Parasuraman, R. (2001). The role of the air traffic controller in future air traffic management: An empirical study of active control versus passive monitoring. *Human Factors*, 43, 519-528.
- Miller, C. (1999). Bridging the Information Transfer Gap: Measuring Goodness of Information Fit. *Journal of Visual Language and Computation*, 10, 523-558.
- Miller, C., Pelican, M. and 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*. Urbana-Champaign, Ill. May.
- Miller, C., Goldman, R., Funk, H., Wu, P. and Pate, B. (2004). 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*. Baltimore, MD; June 7-10.
- Miller, C. and Parasuraman, R. (submitted). "Designing for Flexible Human-Automation Interaction: Playbooks for Supervisory Control." Submitted to *Human Factors*.
- Mosier, K. & Skitka, L. (1996). Human decision makers and automated decision aids: Made for each other? In R. Parasuraman & M. Mouloua, (Eds.), *Automation and human performance: Theory and applications* (pp. 201-220). Mahwah, NJ: Erlbaum.
- Opperman, R., (1994). *Adaptive user support*. Hillsdale, NJ; Erlbaum.
- Parasuraman, R. (1993). Effects of adaptive function allocation on human performance. In D. J. Garland and J. A. Wise (Eds.), *Human factors and advanced aviation technologies* (pp. 147-157). Daytona Beach: Embry-Riddle Aeronautical University Press.
- Parasuraman, R., & Byrne, E. A. (2002). Automation and human performance in aviation., In P. Tsang & M. Vidulich (Eds.) *Principles and practice of aviation psychology*. Mahwah, NJ: Erlbaum, in press.
- Parasuraman, R. Molloy, R. & Singh, I. (1993). Performance consequences of automation-induced 'complacency,' *Int. Jour. of Av. Psych* 3 1-23.
- Parasuraman, R. & Riley, V. (1997). Humans and automation: Use, misuse, disuse, 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.—Part A: Systems and Humans*, 30, 286-297.
- Rasmussen, J. (1986). *Information processing and human-machine interaction*. Amsterdam: North-Holland.

- Rasmussen, J. & Goodstein, P. (1987). Decision support in supervisory control of high-risk industrial systems. *Automatica*, vol. 23, pp. 663-671.
- Riley, V. (1994). *Human use of automation*. Unpublished doctoral dissertation, University of Minnesota.
- Sarter, N. & Woods, D (1994). Pilot interaction with cockpit automation II: An experimental study of pilots' model and awareness of the flight management system. *Int. Jour. of Av. Psych.*, 4, 1-28.
- Sarter, N., Woods, D. & Billings, C. (1997). Automation surprises. In G. Salvendy, (Ed.) *Handbook of human factors and ergonomics*, 2nd ed. (pp.1926-1943). New York:Wiley.
- Satchell, P. (1998). *Innovation and automation*. Aldershot, UK: Ashgate.
- Sheridan, T. (1992). *Telerobotics, automation, and supervisory control*. Cambridge, MA: MIT Press.
- Shneiderman, B. (1997). Direct manipulation for comprehensible, predictable, and controllable user interfaces. *Proceedings of the ACM International Workshop on Intelligent User Interfaces*, New York pp. 33-39.
- Vicente, K. (1999). *Cognitive work analysis: Towards safe, productive, and healthy computer-based work*. Mahwah, NJ; Erlbaum.
- Warm, J. S., Dember, W. N., & Hancock, P. A. (1996). Vigilance and workload in automated systems. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications* (pp. 183–200). Mahwah, NJ: Erlbaum.
- Wickens, C.D. (1992). *Engineering psychology and human performance*. 2nd ed. New York: Harper Collins.
- Wickens, C. (1994). Designing for situation awareness and trust in automation. In *Proceedings of the IFAC Conference*. Baden-Baden, Germany, pp. 174-179.
- Wickens, C., Mavor, A., Parasuraman, R. & McGee, J. (1998). *The future of air traffic control: Human operators and automation*. Washington DC: National Academy Press.
- Weiner, E. (1988). Cockpit automation. In E. Weiner & D. Nagel, (Eds.), *Human factors in aviation* (pp.433-461). San Diego: Academic.
- Wiener, E., & Curry, R. (1980). Flight-deck automation: Promises and problems. *Ergonomics*, 23, 995-1011.