THE INTERNATIONAL JOURNAL OF AVIATION PSYCHOLOGY, *10*(4), 327–341

Automation Management Strategies: Pilot Preferences and Operational Experiences

Wesley A. Olson

Air Command and Staff College
Maxwell Air Force Base, AL

Nadine B. Sarter

Department of Industrial, Welding, and Systems Engineering
The Ohio State University

With the evolution of technology from passive tool to highly independent agent, it is becoming increasingly important to support operators in coordinating human and machine intentions and actions. One way to achieve this goal is the context-sensitive use and implementation of different automation management strategies. This study examined pilots' preferences for and their operational experiences with 3 different strategies—management-by-consent, management-by-exception, and full automation. Pilots expressed a strong preference for management-by-consent in which the automation cannot take action unless explicit pilot consent has been received. High time pressure, high workload, and low task criticality led to a shift in pilots' preferences toward management-by-exception in which the automation can initiate actions on its own. These preferences can be explained, in part, by pilots' operational experiences with existing cockpit systems that illustrate that human–machine coordination is a complex process involving the negotiation of multiple goals, activities, and strategies rather than simply assuming manual control in case of disagreements.

The introduction of highly complex and powerful automation to a variety of domains has led to unexpected difficulties that are the result of an increased need for, but lack of support of, human–machine communication and cooperation. To date,

Requests for reprints should be sent to Nadine B. Sarter, Department of Industrial, Welding, and Systems Engineering, The Ohio State University, 210 Baker Systems Building, 1971 Neil Avenue, Columbus, OH 43210. E-mail: sarter.1@osu.edu

attempts to address these difficulties have focused primarily on either improving the operator's mental model of the system through new forms of training or on developing more effective feedback to improve the communicative skills of the machine. Another possible approach—the context-sensitive choice and implementation of various automation management strategies—has been discussed quite extensively in the literature (e.g., Billings, 1997; Moray, 1986; Sheridan, 1987). Yet, relatively little empirical data exist concerning the feasibility, acceptability, and effectiveness of the different coordination strategies that have been implemented already in a variety of systems and domains.

This article reports on the results of one of the first studies to examine the factors that influence operators' preferences for and their operational experiences with various automation management approaches. This study focuses on the two strategies most frequently implemented in the aviation domain—management-by-consent and management-by-exception. In the former case, the automation is not allowed to take action unless and until explicit operator consent has been received (e.g., the need for the pilot to depress a separate "execute" button before the flight management system [FMS] will complete a pilot-programmed routing change). In contrast, under the management-by-exception approach, the automation can initiate actions independent of the operator who has the option to override or reverse system activities after the fact (Billings, 1997; e.g., automatic transitions between active flight modes).

Earlier research on human–machine collaboration (e.g., Liu, Fuld, & Wickens, 1993; Milewski & Lewis, 1997) showed that the choice between automation management strategies needs to be based on a trade-off decision that considers the expected costs and benefits for different task contexts. In general, system performance can be expected to suffer if the cognitive overhead (Parasuraman & Riley, 1997) associated with the monitoring and negotiation demands of either approach (e.g., tracking the status of unfinished tasks, resolving goal conflicts, planning transitions between tasks) outweighs its potential benefits (e.g., manual workload reduction, increased precision). For example, the required frequent interactions with automated systems under a management-by-consent approach may increase pilots' awareness of, and control over, automation behavior. At the same time, they can lead to difficulties if they occur during already demanding, highly dynamic, high-risk periods—an example of the escalation principle (Woods & Sarter, in press). Management-by-exception, on the other hand, involves fewer negotiations but potentially increases the risk of losing track of system activities. This problem is well documented for modern "glass cockpit" aircraft in which pilots sometimes experience automation surprises when their highly independent systems change their status or behavior on their own (e.g., Sarter, Woods, & Billings, 1997).

To support designers in making informed decisions when selecting and implementing automation management strategies, we need to understand more about the impact of various situational and task factors on the desirability and effectiveness of different approaches. Factors that are likely to play a role include automation properties such as high levels of machine autonomy, authority, complexity, and coupling, which can make it difficult for the human operator to track machine actions, predict future behavior, and intervene when necessary—thus posing a challenge for a management-by-exception approach (Woods, Johannesen, Cook, & Sarter, 1994). This problem is exacerbated in systems that exhibit low observability, that is, those systems that do not actively support operators in monitoring machine status and behavior. Situational and task factors that need to be considered include high levels of workload, time pressure, task difficulty, and task criticality. These factors tend to affect the operator's decision to delegate tasks to a machine, and they may interfere with his or her ability to interact and coordinate with the automation on a frequent basis—a possible requirement of the management-by-consent approach.

The goal of this study is to examine to what extent and in what ways the aforementioned factors influence the desirability and effectiveness of different automation management strategies. To this end, airline pilots were asked to rate and explain their preferences for five different implementations of a possible future flight deck system (one fully automated version with no override option, and two implementations each of the management-by-consent and the management-by-exception approaches) across a wide range of flight scenarios. Because experiences with existing implementations of these strategies are likely to affect pilots' attitudes, and in an effort to capture and learn from these experiences, pilots were also asked to report any difficulties with current flight deck technologies such as the FMS.

Survey data need to be interpreted carefully for a variety of reasons. For example, they are subject to response biases and do not allow for the determination of causality or an objective determination of system performance. We still chose this approach because it represents a very effective and economic tool for the initial exploration of "uncharted territory" based on input from a large number of practitioners (e.g., Sarter & Woods, 1997; Wiener, 1989). This input is valuable in its own right because it tells us about (the reasons for) the likely acceptance or rejection of proposed design solutions and about operational experiences with existing systems (Wickens, 1995). It also serves as the basis for more controlled follow-up performance studies.

METHOD

Participants

The participants in this study were 206 airline pilots flying for two major U.S. air carriers. Our sample included 59 B-757, 84 A-320, and 63 MD-11 pilots. The aircraft types included in the survey—the B-757, A-320, and MD-11—are all modern glass cockpit airplanes that represent different automation philosophies and implementations. A breakdown of the number, age, and experience of the participants is shown in Table 1.

Survey Design

The survey was broken down into three parts. Part 1 provided a detailed description of a future automated air–ground coordination system to be considered by the participants. In Part 2, pilots were asked to rank order the desirability of five different implementations of this system for 15 different datalink–free flight scenarios and to explain the reasons underlying their preferences. Part 3 gathered information about pilots' experiences with existing flight deck systems, most of which represent instantiations of management-by-consent or management-by-exception.

Part 1—system description. In Part 1 of the survey, a possible future datalink–free flight system was described to pilots. Data Link is the envisioned digital air–ground communication medium to be used in future air traffic operations that are expected to allow pilots much more flexibility in choosing and changing their flight path (the so-called Free Flight environment). Although the procedures and equipment for this environment are not yet fully established, it is likely that new forms and higher levels of automation will be required to handle highly flexible flight operations and support effective air–ground coordination in a safe and efficient manner.

Part 2—scenarios. Eight of the factors that were discussed earlier—time pressure, workload, observability, task difficulty, complexity, coupling, task criticality, and autonomy—drove the design of 15 different scenarios. These scenarios provided pilots with specific contexts in which to evaluate the proposed system implementations. The eight variables were manipulated individually within pairs of scenarios, keeping all other factors constant. Table 2 provides summaries of these scenarios along with the variables examined in each case.

TABLE 1
Demographic Statistics of Survey Participants

		Overall Flying Time		*Time in Type*		*Overall Time in Glass Cockpit*		*Age*	
Aircraft Type	*n*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
B-757	59	12,015	5,096	1,465	1,755	2,205[a]	2,038	45.0	5.9
A-320	84	11,536	5,790	1,592	1,614	2,131[b]	1,949	44.1	6.6
MD-11	63	15,150	7,514	1,350	1,327	3,384[c]	2,236	51.5	6.6

[a]Ten B-757 pilots (16.9%) had previously flown the A-320 for an average of 2,037 hr. [b]Eighteen A-320 pilots (21.4%) had previously flown the B-757 for an average of 1,685 hr. [c]Forty-two MD-11 pilots (66.7%) had previously flown the B-757 or B-767 for an average of 3,042 hr.

TABLE 2
List of Scenarios and Associated Variables

Scenario	*Situational Variable*	*Scenario Type*
1A, 1B	Time pressure	Traffic conflict at cruise—immediate versus 30 min in future
1C, 1D	Machine versus human generated resolution	Machine versus human resolution of traffic conflict on takeoff
2A, 2B	Observability	Rerouting during descent with different display features
3A, 3B	Task difficulty	Holding—nonpublished fix 10 miles ahead versus published fix 75 miles ahead
4A, 4B	Task complexity	Four versus two aircraft conflict resolution
5A, 5B	System coupling	The flight management system does or does not have the capability to automatically deploy spoilers to meet request for increased rate of descent
6A, 6B	Task criticality	Radio frequency change versus altitude change
7	System autonomy	Continuous adjustment of airspeed and heading during descent to ensure optimum spacing

Following the description of each scenario, pilots were asked to indicate their preferences by rank ordering five automation options ranging from 1 (*best option*) to 5 (*worst option*).

These five options included one fully automatic option with no capability for pilot override as well as two different variations of both management-by-exception and management-by-consent. The two different variations allowed the pilot to accept or override the instructions either by pressing a single button (called "consent-one" and "exception-one" in the remainder of the manuscript), or by separately accepting or reprogramming different elements of the instruction in the flight management computer, the mode control panel, or both, (referred to as "consent-multiple" and "exception-multi"). Pilots were asked to provide written comments explaining the reasons underlying their preferences in the space provided at the bottom of each page. (See Table 3 for an example scenario and automation options.)

Part 3—experiences with existing automated systems. In this section, using a variation of the critical incident technique (Flanagan, 1954), pilots were asked about their experiences with existing automated systems. Two specific questions were asked: (a) Have you ever experienced a situation in which the automation did less than or more than you expected? and (b) Have you ever experienced the automation to be too difficult or too easy to override? These two questions served to gather information that may explain pilots' preferences expressed in Part 2 of the survey and to identify benefits and disadvantages of specific existing implementations of either management approach.

TABLE 3
Sample Scenario and Automation Options

Scenario 1A

Situation: Cruising at FL-370
Event: Ground–based conflict probe detects immediate traffic conflict (similar to an RA with the current TCAS system) with military fighter aircraft
System: Ground-based ATC computer system calculates and uplinks a conflict resolution

The FMC/MSU System	*The Pilot*	*Preferences for Scenario No. 1A*
Automatically executes uplinked instructions	Cannot override the system	
Automatically executes instructions	Can only override by manually resetting MSU/FMC	
Automatically executes instructions	Can override by depressing “reject” within a reasonable period of time	
Preloads the instructions into the MSU/FMC; does not automatically execute instructions	Executes the entire resolution by selecting “accept”	
Preloads the instructions into the MSU/FMC; does not automatically execute instructions	Executes heading, altitude, and airspeed separately by selecting associated lines on datalink message screen	

Note. FL-370 = flight level 37,000 ft; RA = resolution advisory; TCAS = traffic alert and collision avoidance system; ATC = air traffic control; FMC = flight management computer; MSU = mode selection unit.

Procedure

After obtaining airline and union permission, the survey was distributed via company mail to all pilots currently flying the designated aircraft for each airline. A cover letter accompanying each survey explained the general purpose of the survey—to better understand and provide a basis for improving coordination between pilots and automated systems. The letter stressed the confidentiality of all responses and stated that this survey provided an opportunity to influence the design of future technology. Participation was voluntary, and participants were not reimbursed for their efforts. Surveys were initially distributed to 750 B-757 pilots, 750 A-320 pilots, and 350 MD-11 pilots. The relatively low return rate (11%; 206 returned surveys out of 1,850 distributed) may be explained by the considerable effort required to complete this survey.

RESULTS

Pilot Ratings of Different Automation Management Strategies

Figure 1 shows the median ratings for the five different automation options for all scenarios. The data are collapsed across aircraft types and experience levels because no significant differences were found between the responses of B-757, A-320, and MD-11 pilots or between pilots at different levels of flight experience.

Multiple comparisons using the Friedman two-way analysis of variance by ranks confirmed the general pattern emerging from these data. The fully automated option received the worst rating for all scenarios, followed by the two management-by-exception options—multiple button reject first, and next, the single reject button option. The two management-by-consent options were rated best, with one button consent being preferred for Scenarios 1A (high time pressure; $p < .05$), 1C (machine-generated conflict resolution advisory; $p < .05$), 4A and 4B (high vs. low task complexity; $p < .05$), and 6A and 6B (high vs. low task criticality; $p < .05$).

Because median data are not particularly diagnostic and can mask underlying response patterns, the best option and worst option frequencies were examined. Figure 2 shows the frequency counts for the best option ratings across scenarios.

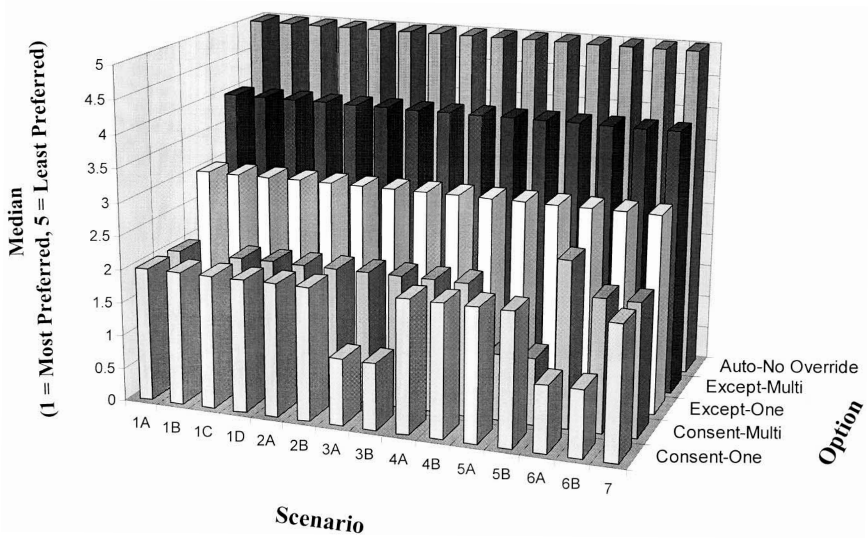

FIGURE 1 Median ratings of automation options across all scenarios.

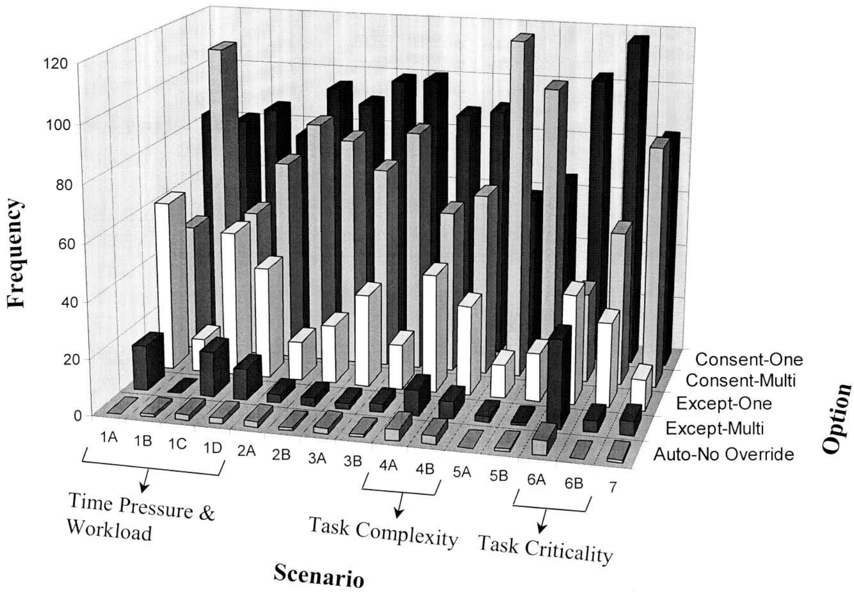

FIGURE 2 Frequency of best ratings for each automation option across all scenarios.

The pattern of best option frequencies mirrors the median rating results. Note, however, that in two of the scenarios—1A (the immediate traffic conflict) and 6A (the radio frequency change)—pilots not only preferred one button consent to multiple button consent; they also ranked the one button exception option as the best option compared to the multiple button consent option. Again, the fully automatic option was overwhelmingly cited as the worst option in all scenarios. It accounts for over 90% of all "worst" rankings in every scenario.

We reviewed the pilots' written comments to identify common reasons underlying pilots' preferences. Desire for control was the reason most often cited for selecting management-by-consent (46.9%), whereas the converse, a perceived lack of control, led most frequently to the rejection of both management-by-exception options (62.6%) and the fully automatic option (74.4%). Speed of response was cited as the primary reason for preferring both management-by-exception (36.6%) and fully automatic without override (21.4%) approaches, whereas lack of time for a human response (62.5%) and workload considerations (25.0%) were cited as major reasons for rejecting the management-by-consent options. Finally, low task criticality was another reason for preferring both the fully automatic option (7.1%) and the two management-by-exception options (4.1%).

Operational Experiences With Various Automation Management Strategies

Because prior exposure to existing implementations of the two automation management approaches may have influenced pilots' attitudes and preferences, they were asked about their experiences with existing flight deck technology. The first question was, "Have you experienced cases in which automation did too much or too little?" Pilots (151 out of 193; 78.2%) responding to this question reported that they had been surprised by the automation. A further breakdown reveals an almost even split between experiences in which the automation did more than expected (39 cases), less than expected (55 cases), or both (57 cases).

Pilots were also asked about their experiences with overriding automated systems. Overall, 43.9% of the pilots (83 out of the 189 pilots who responded) considered the effort that is required to override the automation to be inadequate. Sixty-five pilots indicated that the automation was too difficult to override. Fifteen pilots reported they had experienced automation being both too easy and too difficult, and only 3 pilots felt that it was too easy to override.

The examples that pilots provided in response to both questions overlapped considerably and were therefore combined for the purposes of analysis and categorization (see Table 4). It is impossible within the scope of this article to discuss in detail the nature of reported problems. It is clear, however, that the majority of difficulties are related to reinstructing automated systems and to the various ways in which machine intentions and behavior can conflict with those of the human operator (i.e., wrong style, wrong goal, coupling problems, or incorrect priorities).

DISCUSSION

With the evolution of modern technology from reactive tool to powerful and independent agent, it is becoming increasingly important to support operators in effectively tracking and coordinating the intentions and actions of the joint human–machine team. Efforts are ongoing to achieve this goal through measures such as improved feedback design and new forms of operator training. Yet, another approach, the context-sensitive choice and implementation of automation management strategies, has received less attention to date. Although the topic has been discussed extensively in the literature, relatively little empirical evidence exists concerning the desirability and effectiveness of different management strategies and implementations for various task contexts. This situation poses a challenge for designers who, despite this lack of data, cannot avoid making decisions about automation management approaches when developing mod-

TABLE 4
Frequency of Top Nine Reported Difficulties With Existing Automated Systems

Type of Difficulty	*Frequency*	*%*
Automation over- or underreacted	28	15.5
Automation difficult to instruct	27	14.9
Automation difficult to reinstruct	27	14.9
Automation pursued wrong goal	26	14.4
Automation performance difficult to compare to pilot expectations	19	10.5
Automation pursued correct goal but did more than expected	9	5.0
Automation performed some (but not all) expected actions	8	4.4
Automation applied incorrect priorities	8	4.4
Effects and indications of automation actions difficult to observe	8	4.4

ern technology. In an attempt to provide input to their design decisions and to further our understanding of human–machine coordination, this survey gathered both subjective and experiential data on automation management approaches.

Perceived control of automated systems turned out to be very important to pilots in this survey and explained their overwhelming preference for the management-by-consent approach for all scenarios except those involving high levels of time pressure and workload or a low level of task criticality. For situations involving one of those three factors, a considerable number (but still not a majority) of pilots preferred management-by-exception instead.

Pilots' ratings need to be interpreted carefully, however, because they can be explained, in part, by their experiences with existing technology. In particular, their difficulties with overriding current automated systems may have resulted in their rejection of the management-by-exception approach, which requires that midcourse corrections be made quickly and effectively. Thus, a different implementation of this strategy may well lead to increased acceptability. Pilots' experiences with current systems also illustrate that quick and effective intervention is critical not only in case of goal conflicts, but also in situations that are characterized by conflicting styles and priorities.

The Importance of (Perceived) Control

Pilots' responses appear to support the human-centered design principle (Billings, 1997) that operators should have ultimate control over automated systems. This is indicated by their strong dislike for the fully automatic option with no override capability, and it explains, in part, their preference for management-by-consent over management-by-exception. Although both management-by-consent and management-by-exception allow pilots to intervene with system actions, some existing systems make it very difficult and effortful for pilots to redirect intentions and actions. Thus, the rejection of management-by-exception can be explained by its current implementations, which tend to provide nominal rather than effective control of the automation.

It is also important to note that although pilots expressed a strong dislike for the proposed automation without override capability, they did not report negative experiences with such systems that already exist on many advanced transport aircraft. There are several possible explanations for this apparent paradox. The absence of comments may simply reflect a lack of experience; some of these systems become active only in a small number of critical situations. Other systems may have led to positive experiences due to their high degree of speed, effectiveness, and consistency in performing highly critical tasks such as wind shear recovery. Finally, some fully automatic systems may be accepted because they perform tasks that were previously accomplished by the flight engineer (e.g., MD-11 automatic subsystems controllers). Follow-up studies to further examine this issue may provide useful insights into the successful use and implementation of highly independent and authoritative technology.

The Effects of Time Pressure, Workload, and Task Criticality

Pilots' preference for a particular implementation of management-by-consent was affected by the nature of the specific scenario under consideration. Both the median data and the frequency data show that, for situations involving high time pressure, high task complexity, or low task criticality, pilots prefer the option that entails a single "accept" button to the one requiring them to separately accept or reject individual parts of a clearance. This may, at first, seem to contradict pilots' strong desire for control that would appear greater if they can make decisions about each individual component of a clearance. However, their ratings appear to acknowledge the fact that such increased control comes at a price—it brings with it an increase in cognitive demands and required interaction with the system. Thus, pilots' preferences reflect a trade-off between minimizing costs (such as attentional and interface management demands) and retaining benefits (such as control) to the extent possible.

An analysis of pilots' written comments along with a closer examination of the frequency data show that the aforementioned variables, high time pressure and workload, high task complexity, and low task criticality, also lead a considerable number (although not a majority) of pilots to prefer the management-by-exception option for Scenarios 1A, 1C, and 1D. The most often cited reason for selecting this option is that it affords a fast response (as required in the case of the immediate traffic conflicts in Scenarios 1A, 1C, and 1D) while still allowing the pilot to override the automation if necessary.

Observed Difficulties With Directing and Redirecting Automated Systems

In an attempt to capture and learn from operational experience, we asked pilots about difficulties they may have experienced with existing flight deck systems that represent instantiations of the different automation management strategies.

Of all pilots in this survey, 78.2% indicated that they had experienced breakdowns in human–automation coordination and resulting automation surprises. This result replicates similar findings by Wiener (1989) and by Sarter and Woods (1992, 1994, 1997) and extends them to yet another glass cockpit aircraft, the MD-11.

The nature of the reported coordination breakdowns has implications for the feasibility and acceptability of the management-by-consent and management-by-exception approaches. Under a management-by-exception approach, intervention is required only when the pilot disagrees with machine actions; under a management-by-consent approach, however, pilot input is always required before any action may take place. Although this may suggest an advantage for the management-by-exception approach, this advantage will only materialize if the frequency of pilot–machine disagreements and, thus, the associated negotiation demands are very low. However, pilots in this survey reported the opposite—disagreements are relatively frequent. The result is new attentional and interface management demands and, possibly, mistrust in the automation (Lee & Moray, 1994; Parasuraman, Molloy, & Singh, 1993).

As mentioned earlier, another factor that may explain pilots' rejection of the management-by-exception approach is reported difficulties with overriding existing automated systems. Approximately 75% of the override difficulties described in this survey involved problems with reprogramming the automation to fine-tune its behavior as opposed to reverting to manual control. This finding supports earlier criticism of the ubiquitous call for pilots to turn off the automation when it does not act as expected or desired. Pilot–automation collaboration is a far more complex process in which the human and machine agent try to coordinate and negotiate multiple intentions, actions, and strategies rather than eliminating one another completely in response to a disagreement about one single aspect of their performance.

Coordination—More Than Resolving Goal Conflicts

Coordination can be defined as the "management of interdependencies" (Malone & Crowston, 1990, p. 358). Although goal conflicts are one important and often discussed type of interdependency, many others need to be considered. Our data indicate that breakdowns in human–machine coordination with current systems tend to occur most often when the automation (a) pursues unwanted goals or targets (e.g., flies at an undesired speed); (b) pursues the correct goals, but sets wrong priorities or pursues them in an unintended manner (e.g., gives smooth flight priority over meeting bottom of descent constraints or moves throttles too abruptly); and (c) partially pursues the correct goal but does more or less than pilots expect (e.g., on some aircraft, inserting a new takeoff runway can delete the entire route of flight). Although current systems allow pilots to change most machine goals, they often fail to support modifications or adjustments of system priorities and styles. A lack of control over these important parameters can lead pilots to adopt unanticipated and sometimes maladaptive system and task tailoring strategies (e.g., Cook, Woods, McColligan, & Howie, 1991).

Future Research Needs

This survey represents a first step toward a better understanding of the impact of various contextual factors on the desirability and effectiveness of different automation management strategies and implementations. It also suggests a number of important questions that need to be addressed in future research.

First, although pilots overwhelmingly preferred a management-by-consent approach, they reported relatively few difficulties with existing highly automatic systems that do not allow for pilot control or intervention. Further research is needed to better understand when the use of highly independent and authoritative technology is acceptable or even desirable.

Given the often observed considerable differences between preference and performance data (Andre & Wickens, 1995), the findings of this survey must be validated by empirical evaluations of joint system performance across a wide range of tasks and contexts. With this goal in mind, we have recently conducted a first follow-up simulator study to explore actual (as opposed to assumed) performance with the widely preferred management-by-consent approach. In particular, the study examined the effects of conflict type, time pressure, display design, and trust on pilots' ability to track and anticipate the goals, priorities, and styles of an automated flight deck system. Maintaining awareness of all of these aspects of system performance is critical for giving informed (as opposed to perfunctory) consent to proposed machine actions (Olson & Sarter, 2000).

Because different tasks and situations are likely to require transitions between automation management strategies, research must explore ways to effectively support responsibility awareness, that is, the knowledge of roles and responsibilities of each entity at all times (Coury & Semmel, 1996). This is critical to avoid situations in which shared control leads to responsibility diffusion, which, in turn, can lead to delayed or missing interventions by the human operator or to a struggle between human and machine (e.g., Darley & Latane, 1968; Sarter & Woods, 1995).

Finally, more effective protocols for human–automation negotiation need to be developed to better support rapid and effective interventions in the context of a management-by-exception approach. This will be particularly important for environments such as the envisioned air traffic management system in which a large number of human and machine agents in different locations need to collaborate to ensure the safety of highly dynamic traffic operations.

ACKNOWLEDGMENTS

The preparation of this article was supported, in part, by research Grant 96–G–043 from the Federal Aviation Administration (with Tom McCloy and Eleana Edens as technical monitors).

We appreciate the cooperation of the airlines and pilots who agreed to participate in this survey.

REFERENCES

Andre, A. D., & Wickens, C. D. (1995). When users want what's not best for them. *Ergonomics in Design, 3*(4), 10–14.

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

Cook, R. I., Woods, D. D., McColligan, E., & Howie, M. B. (1991). Cognitive consequences of "clumsy" automation on high workload, high consequence human performance. In R. T. Savely (Ed.), *Fourth Annual Workshop on Space Operations, Applications and Research, SOAR '90* (NASA Rep. No. CP–3103). Washington, DC: NASA.

Coury, B. G., & Semmel, R. D. (1996). Supervisory control and the design of intelligent interfaces. In R. Parasuraman & M. Mouloua (Eds.), *Automation and human performance: Theory and applications* (pp. 221–242). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Darley, J. M., & Latane, B. (1968). Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology, 8,* 377–383.

Flanagan, J. C. (1954). The critical incident technique. *Psychological Bulletin, 51,* 327–358.

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

Liu, Y., Fuld, R., & Wickens, C. D. (1993). Monitoring behavior in manual and automated scheduling systems. *International Journal of Man–Machine Studies, 39,* 1015–1029.

Malone, T., & Crowston, K. (1990). What is coordination theory and how can it help design cooperative work systems? *Proceedings of the ACM Conference on Computer-Supported Cooperative Work, USA,* 357–370.

Milewski, A. E., & Lewis, S. H. (1997). Delegating to software agents. *International Journal of Human–Computer Studies, 46,* 485–500.

Moray, N. (1986). Monitoring behavior and supervisory control. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance* (Vol. 2, pp. 40-1–40-46). New York: Wiley.

Olson, W. A., & Sarter, N. B. (2000). *Management-by-consent in human–machine systems: When and why it breaks down.* Manuscript submitted for publication.

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

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

Sarter N. B., & Woods, D. D. (1992). Pilot interaction with cockpit automation: Operational experiences with the flight management system. *The International Journal of Aviation Psychology, 2,* 303–321.

Sarter N. B., & Woods, D. D. (1994). Pilot interaction with cockpit automation II: An experimental study of pilots’ model and awareness of the flight management system. *International Journal of Aviation Psychology, 4,* 1–28.

Sarter N. B., & 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.

Sarter, N. B., & Woods, D. D. (1997). Teamplay with a powerful and independent agent: A corpus of operational experiences and automation surprises on the airbus A-320. *Human Factors, 39,* 553–569.

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

Sheridan, T. B. (1987). Supervisory control. In G. Salvendy (Ed.), *Handbook of human factors* (pp. 1243–1268). New York: Wiley.

Wickens, C. D. (1995). Aerospace techniques. In J. Weimer (Ed.), *Research techniques in human engineering* (pp. 112–140). Englewood Cliffs, NJ: Prentice Hall.

Wiener, E. L. (1989) *Human factors of advanced technology (“glass cockpit”) transport aircraft* (NASA Contractor Rep. No. 117528). Moffett Field, CA: NASA Ames Research Center.

Woods, D. D., Johannesen, L. J., Cook, R. I., & Sarter, N. B. (1994). *Behind human error: Cognitive systems, computers, and hindsight* (CSERIAC State-of-the-art Rep., SOAR ’94–01). Wright-Patterson Air Force Base, OH: CSERIAC.

Woods D. D., & Sarter, N. B. (in press). Learning from automation surprises and “going sour” accidents. In N. B. Sarter & R. Amalberti (Eds.), *Cognitive engineering in the aviation domain.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Manuscript first received February 1999