Human Factors: The Journal of the Human Factors and Ergonomics Society
http://hfs.sagepub.com/

Toward a Characterization of Adaptive Systems : A Framework for Researchers and System Designers

Karen M. Feigh, Michael C. Dorneich and Caroline C. Hayes

Human Factors: The Journal of the Human Factors and Ergonomics Society 2012 54: 1008 originally published online 1 May 2012
DOI: 10.1177/0018720812443983

Toward a Characterization of Adaptive Systems: A Framework for Researchers and System Designers

Karen M. Feigh, Georgia Institute of Technology, Atlanta, Georgia, **Michael C. Dorneich**, Honeywell Laboratories, Golden Valley, Minnesota, and **Caroline C. Hayes**, University of Minnesota, Minneapolis, Minnesota

Objective: This article presents a systematic framework characterizing adaptive systems.

Background: Adaptive systems are those that can appropriately modify their behavior to fit the current context. This concept is appealing because it offers the possibility of creating computer assistants that behave like good human assistants who can provide what is needed without being asked. However, the majority of adaptive systems have been experimental rather than practical because of the technical challenges in accurately perceiving and interpreting users' current cognitive state; integrating cognitive state, environment, and task information; and using it to predict users' current needs. The authors anticipate that recent developments in neurological and physiological sensors to identify users' cognitive state will increase interest in adaptive systems research and practice over the next few years.

Method: To inform future efforts in adaptive systems, this work provides an organizing framework for characterizing adaptive systems, identifying considerations and implications, and suggesting future research issues.

Results: A two-part framework is presented that (a) categorizes ways in which adaptive systems can modify their behavior and (b) characterizes trigger mechanisms through which adaptive systems can sense the current situation and decide how to adapt.

Conclusion: The framework provided in this article provides a tool for organizing and informing past, present, and future research and development efforts in adaptive systems.

Keywords: adaptive systems, adaptive automation, adaptable automation, dynamic function allocation

Address correspondence to Karen M. Feigh, Georgia Institute of Technology, 270 Ferst Dr., Atlanta, GA 30332-0150, USA; karen.feigh@gatech.edu.

HUMAN FACTORS
Vol. 54, No. 6, December 2012, pp. 1008-1024
DOI:10.1177/0018720812443983

INTRODUCTION

Adaptive systems are the technological component of joint human–machine systems that can change their behavior to meet the changing needs of their users, often without explicit instructions from their users. Adaptive systems do so by tracking and sensing information about their users, their current tasks, and their environment. Our motivation for examining adaptive systems at this time is that recent advances in sensor technology have created a renewed focus on adaptive systems among both researchers and software developers. The goal of our work is to assist this new surge of adaptive system researchers and software developers by providing a high-level characterization of adaptive systems and adaptive systems research.

The contribution of our framework is that it expands existing perspectives on what adaptive systems are. Previous characterizations of adaptive systems tended to focus primarily on one type of adaptation: modifications of the allocation of functions performed by the user and the adaptive system. However, there are many other ways in which adaptive systems can adapt; for example, by changing the amount of detail presented to users (which can be important when they are very stressed) or by changing the sensory modality in which information is presented: visual versus auditory. Our framework not only captures many *types of adaptations* but also describes and categorizes a variety of the *trigger mechanisms* by which the adaptations are invoked or disengaged. For each category, we provide technical descriptions, review the implications and considerations, and provide concrete examples drawn from the literature. We also discuss past and current challenges in creating successful adaptive systems, recent work to overcome those challenges, and future research directions to make adaptive systems more practical, usable, and reliable.

We start by describing adaptive systems in more detail. Adaptive systems may be components of larger joint human–machine systems such as the autopilot, warning, or navigation systems that assist a pilot in flying an airplane, a tutoring system that trains students, or a help system on a mobile device. The potential advantage in making such systems adaptive is that users' needs may change as their current task, environment, and cognitive state also change. *Cognitive state* refers to properties such as the user's mental workload, fatigue, or stress. For example, if a user is driving slowly through an urban environment, an adaptive GPS system might automatically adjust the map to show more detail than when the user is driving at high speed on a relatively featureless desert highway. As a second example, if an adaptive system senses that a user is focused on managing a complex and dangerous emergency, such as landing an airplane after an engine catches fire, it may prevent all interruptions except for those related to the most essential and pertinent alerts (Dorneich, Mathan, Whitlow, & Ververs, 2010). When implemented appropriately, adaptive systems act as skilled human assistants that unobtrusively observe their supervisors' actions and state of mind, comprehend the evolving situation, and provide appropriate assistance without being asked. Such systems have the potential to enhance joint human–machine system performance.

To illustrate the general structure of an adaptive system, Figure 1 depicts a generalized flow diagram, patterned after the "perceive, select, act" cycle through which intelligent agents process information, make decisions, and interact with the environment (Wickens, 1992). Data collected via sensors or data feeds from various sources can be used to make assessments of the state of the system, the external environment, the task, and the user. The assessments can then be used, in isolation or combination, to trigger an adaptations manager's decision of which adaptations to select. The adaptive changes are executed by the automation and the human–machine interface. For example, the context assessment module may collect physiological sensor information that is classified into levels of workload. In addition to a workload assessment, the user's current tasks and the system's current state are all used by the adaptations manager to decide that high workload might compromise the user's performance on the highest priority task; it offloads lower-priority tasks to automation so the user can focus on only the most critical task. Once this task is finished, the adaptations manger then decides to return to regular operation and disengage the adaptation.

Adaptive systems have remained limited because of the difficulty in assessing context; for example, they have often relied on static task models and user performance to gauge user state indirectly. However, recent advances in real-time, noninvasive, user cognitive state assessment have opened up the possibilities for more sophisticated adaptive systems that can sense user state directly. Nonetheless, multiple technological challenges remain, including the need for (a) more robust, accurate, wearable, and unobtrusive neurological and physiological sensors capable of providing the real-time information needed to determine user's cognitive state; and (b) a better understanding of how to translate information on the user's cognitive state, task, and environment into meaningful guidance for an adaptive system. These and other challenges made it difficult for adaptive systems to correctly and accurately ascertain the current situation and needs (Bainbridge, 1983). When an adaptive system fails to do so, it may provide wrong, inappropriate, or untimely support, which can reduce the user's effectiveness, become an annoyance, and, in the worst case, compromise safety. For example, an adaptive system that provides nonessential weather updates instead of nearby emergency landing sites to a pilot during an engine failure is not only a distraction and an annoyance but also a safety concern. An adaptive system that is not sufficiently accurate at assessing current needs is like an unobservant or poorly trained assistant who provides more hindrances than help (Miller & Funk, 1997; Miller & Hannen, 1999).

Although there were successful adaptive systems 10 years ago, they did not often use information about the user's cognitive state to accomplish their jobs. For example, computer

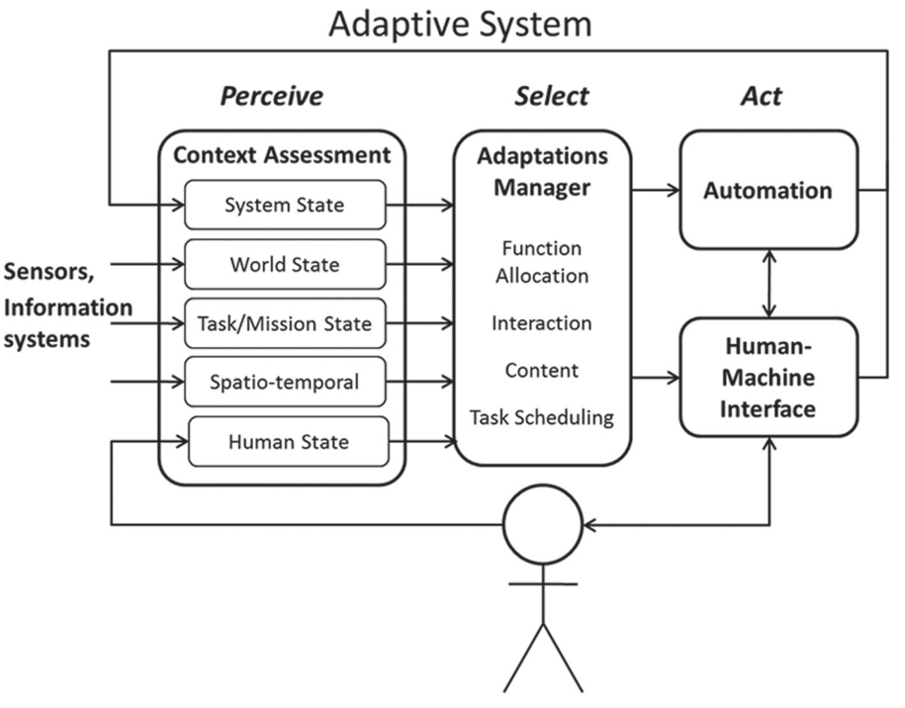

Figure 1. Diagram for a generic adaptive joint human–machine system.

tutors could adapt their teaching style as the knowledge of the student grew (Dorneich & Jones, 2001; Johnson, Shaw, Marshall, & LaBore, 2003). However, such systems were limited in their ability to determine the user's state of frustration or satisfaction and thus often added to their frustration. At that time, it was very difficult to use brain and physiological sensors to provide an adaptive system with real-time feedback on the user's cognitive state because the sensors were so cumbersome to wear (e.g., many sensors, glued to the head, wired to computers), rendering them completely impractical for field and mobile applications. Signals from an electroencephalogram (EEG) were hard to read accurately, particularly in noisy environments. Methods for translating signals into cognitive state information must constantly be tuned and retuned for each individual. These are only a few of the many challenges.

Recent advances in multiple disciplines have brought adaptive systems that use cognitive state information closer to the practical realm. Examples of sensor advances include greater sensitivity, enabling them to more accurately detect very weak signals from the brain, even in noisy environments (Mazaeva, Dorneich, Mathan, & Ververs, 2005); and reductions in size, weight, and power requirements, rendering them easier and more practical to wear. These advances have enabled a new generation of "user sensitive" adaptive systems to be successfully built, fielded, and tested (Dorneich, Ververs, Mathan, Whitlow, & Reusser, 2006; Prinzel, Freeman, Scerbo, Mikulka, & Pope, 2000; Scerbo et al., 2001). This has enabled the research community to experimentally measure some of the benefits that can be gained from adaptive systems and to identify some of the costs and practical complexities of using them. What has emerged is a richer understanding of the practical implications of using adaptive systems as intelligent assistants.

In this work, we aim to convey this emerging picture by first presenting a two-part framework characterizing the following:

- *Adaptation types* describe ways in which a system's behavior or interface can be adapted; for example, temporarily take over tasks for users to

ease their burden during busy or stressful times, change the type and detail of information presented, or change the sensory mode (e.g., system, visual, auditory, haptic) in which users exchange information with the system.
- *Adaptation triggers and methods* are used to decide when and how to change the system's behaviors or properties.

In addition, we discuss areas of future research needed to make adaptive systems more practical, usable, and beneficial, particularly those that use sensor input from users.

The framework extends existing integrative work characterizing adaptive systems, which has primarily focused on adaptations of function allocation, in which the system takes over one or more of the tasks performed by the human to lighten the user's workload (Endsley, 1987; Endsley & Kaber, 1999; Parasuraman, Sheridan, & Wickens, 2000; Sheridan & Verplank, 1978). Our framework goes beyond this to also describe other types of adaptations, including changes to the interface content, interaction, or task management. In addition, we characterize the triggering methods that include additional structure and effort necessary to manage a dynamically adaptive system on the part of both the system developer and the user. Our framework follows in the tradition of Rouse (1988), who postulated early on a broad framework for use in the conceptual design of what was then called adaptive aiding, and Rothrock, Koubek, Fuchs, Haas, and Salvendy (2002), who created a three-dimensional taxonomy to describe the effectiveness of an interface for task execution. The framework takes a broader view than Rothrock et al. and provides more detail than Rouse.

The intended audience includes designers, evaluators, and researchers of adaptive systems. The goals behind the framework are to (a) assist designers and evaluators of adaptive systems to systematically consider a broader range of system adaptations, methods for triggering such adaptations, and the trade-offs that may result from specific design choices, (b) tie together a diverse body of literature on adaptive systems, and (c) provide a context in which future research can be situated. The framework description is not exhaustive, nor is that the goal. Instead, the article aims to provide a useful guide for system designers, evaluators, and researchers.

The following sections describe the Taxonomy of Adaptations and the Taxonomy of Triggers and discuss their design considerations, implications, common benefits, and caveats. Where applicable, the literature is used to provide examples of the concepts presented. Supported by the framework, the article concludes by outlining areas where future research is needed.

TAXONOMY OF ADAPTATIONS

Figure 2 illustrates the Taxonomy of Adaptations for human–machine systems. The top-level categories capture the gamut of possible adaptations and show the four primary ways in which a system developer might make the automated portion of a human–machine system adaptive so that it may better meet the needs of the current situation:

- *Modification of function allocation*. One can dynamically change *who* (human or machine) performs each function, task, or subtask. For example, the third generation of the Traffic Collision Avoidance System (TCAS III) will take over the function of flying an airplane when a collision with another aircraft is imminent; human reflexes are sometimes too slow to avoid the crash if the other plane is very close when it is first detected (Botargues, 2008).
- *Modification of task scheduling*. A system may be designed to dynamically change when tasks are performed, including their duration and priority. For example, some smart phones change the ring tone to silent or vibrate when the calendar on the phone indicates that the operator has a scheduled activity.
- *Modification of interaction*. A system may be designed to dynamically change how it interacts with the users. Examples include changing the layout of a visual interface, the mode in which information is presented and received (e.g., visual, auditory, haptic), whether information is exchanged synchronously or asynchronously, and whether information is pushed or pulled. For example, the Communications Scheduler (Dorneich et al., 2010) adapts soldiers' communications during

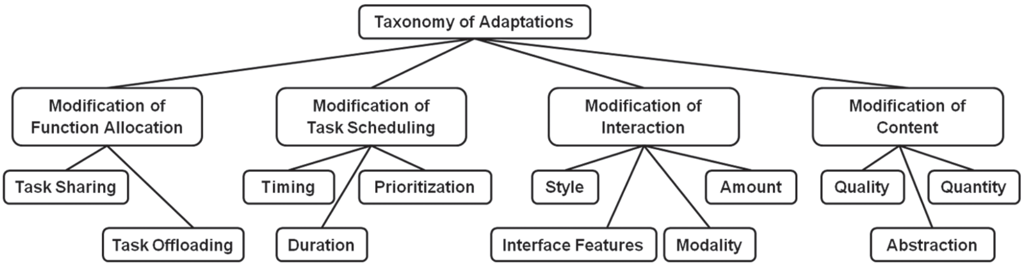

Figure 2. Taxonomy of Adaptations for adaptive systems.

high workload times and changes the interaction so that instead of pushing information to soldiers as it comes in, the soldier must pull information from the system when workload allows.

- *Modification of content*. A system may be designed to dynamically change what information it presents to the user, including what categories of information are presented and at what level of detail or abstraction. For example, a map display for a car may sense its GPS position and speed, and can automatically adjust the scale of the map accordingly, changing the information content by providing detailed information when traveling slowly though an urban area and a larger view when traveling at highway speeds through a rural area.

Although the top-level categories are mutually exclusive, adaptations in one category may often be accompanied by adaptations in another. It is likely that a particular adaptive system will include adaptations in more than one of the four primary categories. For example, the Communications Scheduler (CoS) described above contains adaptations in three categories: function allocation, task scheduling, and interaction. By sensing when the user is experiencing high workload, it invokes a change in function allocation to take over the function of message triage. The CoS further changes the interaction from pushing data to requiring the user to pull data and influences task scheduling by modifying the order in which messages are displayed. The following sections discuss each type of modification in detail, including their implications and considerations, and provide examples.

Who: Modification of Function Allocation

Modification of function allocation is the process of dividing functions (or tasks) between people and machines and deciding who (or what) should perform each task. Function allocation is often considered the same concept as adaptive automation. This article uses adaptive automation as a broader descriptor that includes function allocation as well as the other three categories in the Taxonomy of Adaptations. If system designers allocate functions between humans and machines when they design the system, it is referred to as *static function allocation*. If they design the system so that tasks can be reassigned while the system is in use, it is referred to as *dynamic function allocation* (DFA; or adaptive function allocation). The latter is the primary concern in this work. A number of articles provide comprehensive reviews of systems that adapt themselves through DFA (Kaber & Endsley, 2004; Scerbo, 1996). Changes in function allocation consist of changes to the *assignment* of functions to each agent and have significant implications for the distribution of *responsibility* and *authority* among agents. Taking an operator-centric view, function allocation manifests itself in a combination of *task sharing* (between the operator and automation) and *task offloading* (from the operator to the automation).

Assignment indicates which agent has been assigned to perform a specific function or task. The choice of task assignment is not simple, and many function allocation guidelines and corresponding critiques have surfaced over the years (Abbott & Rogers, 1993; Alter et al., 1995; Dekker & Woods, 2002; Duncan, 1986; Fitts, 1951; Jordan, 1963; Lintern, 2012; Price, 1985). Out of this literature, several function allocation categorizations have been developed to describe varying degrees of automation called levels of automation (LOAs). Initial LOAs assumed that the degree of automation varied along a unidimensional continuum; at the lower end the human performs all tasks, at the upper end the machine performs all tasks (Endsley, 1987; Endsley & Kaber, 1999; Sheridan & Verplank, 1978). In the middle, only portions of a task are automated, creating a third category of "shared performance" (McGuire et al., 1991; Meister, 1985; Sheridan, 2000; Tenney, Rogers, & Pew, 1995).

LOAs have been refined through the addition of a second dimension corresponding to specific information processing stages (Parasuraman et al., 2000). This decomposition has been extensively used in the adaptive automation domain, by varying the LOA to improve performance (Kaber & Endsley, 2004; Kaber, Wright, Prinzel, & Clamann, 2005; Kaber, Wright, & Sheik-Nainar, 2006).

Implications and considerations. Changes to function assignment have significant implications for the distribution of responsibility and authority between agents.

Responsibility. Responsibility indicates which agent is responsible for the outcome of a specific function or task. The responsible agent is not always the same as the assigned agent, especially in the case of joint human–machine systems, which include both human and automated agents. Historically, automated agents have been immune from responsibility, but the drive to minimize human error is changing this as automation is given more responsibility. Dividing the functionality between two (or more) agents requires particular care because dividing the work has an effect beyond just shifting task assignment between agents. Research has shown that the allocation of partial functionality to automation actually changes the nature of the work for the operator because the assignment and responsibility have been split (Billings, 1997; Miller & Parasuraman, 2007; Parasuraman et al., 2000; Parasuraman & Riley, 1997; Woods, 1996). The split between assignment and responsibility adds additional "induced" functions such as monitoring, communicating, and coordinating (Lee, Kim, & Feigh, 2009).

Authority. Authority indicates the level of control an agent has to modify the execution of the tasks and functions to achieve a goal, including changes to assignment. Authority is especially critical in off-nominal situations. Authority differs from responsibility in that authority affects the manner in which a goal or outcome is achieved, whereas responsibility affects the actual performance of the system toward that goal. Because the authorizing agent may mandate the way in which the responsible agents may act, it is possible to limit the ability of the responsible agent to meet its obligations, resulting in an authority–responsibility double bind (Woods, 1996).

As changes to function allocation are made to task assignment, the use of LOAs to describe these changes has proven very beneficial to system designers to easily communicate concepts of operation for new automated systems. However, the reliance on LOAs to describe function allocation has significant implications because LOAs do not capture differences imposed on responsibility and authority. Miller and Parasuraman (2007) argue that, for use with DFA, LOA decomposition needs to be extended beyond the four information processing stages and suggest a delegation method. They assert that tasks are routinely accomplished by hierarchical, decomposable sequences of activities, and it is necessary to differentially apply automation to every subtask. Over time, they believe that multiple, alternate decompositions will be needed depending on the context, where each alternative has a different combination of human and automation subtasks and consequently uses different methods to accomplish the parent task. Accordingly, a single adaptive system may occupy *multiple* points on the LOA continuum, and some systems may be impossible to classify using an LOA.

DFA imposes additional requirements for verification and validation of automated systems because consideration of overall human–computer performance in an adaptive system may come at the expense of suboptimization of local and specific task performance. In fact, this is a primary tenet of a human-centered design philosophy (Norman, 1986; Palmer, 1995). Accordingly, although changes to function allocation may enhance, preserve, or degrade capabilities at the task level, the joint human–machine system should be improved in some way. Usually this improvement comes through the reduction or leveling of workload, increase of situation awareness, and improved robustness to unforeseen and nonnominal circumstances. Wickens, Li, Santamaria, Sebok, and Sarter (2010) used a meta-analysis to examine the degree of automation in both nominal and off-nominal conditions and found that the increase in performance with higher automation is accompanied by an increase in costs for fallible automation but that these results are mediated by situation awareness.

Other considerations surround the task transition between humans and automation; the transition requires explicit coordination and will create additional management and communication work for both agents. Finally, loss of skill, knowledge, and situation awareness may result from too frequent use of automation. It is well documented that use of automation can result in decay of skill or knowledge (Byrne & Parasuraman, 1996; Hancock, Chignell, & Loewenthal, 1985; Parasuraman & Bowers, 1987). Thus, adaptive systems, which automate some tasks only when needed, may help to mitigate these concerns as they provide opportunities for the human to conduct the task whenever possible.

Examples. Kaber and colleagues (Kaber & Endsley, 2004; Kaber et al., 2005; Kaber et al., 2006) describe a series of investigations into adaptive systems that change the function allocation dynamically using an LOA to define the function allocations. According to Kaber and Endsley (2004), LOA was the dominant factor affecting performance. Unfortunately, the "best" LOA combination depended on the role and metric assessed. In addition, Arciszewski, de Greef, and van Delft (2009) outlined an adaptive system that transitions between an automated and manual mode for target classification. The system adaptively changes modes when it has difficulty identifying the target. The use of the automated mode offloads routine target classification tasks and allows the operator to focus on more difficult cases.

When: Modification of Task Scheduling

The modification of the task scheduling category of the taxonomy describes automation adaptations to support individuals in multitasking, interruption-laden environments. Task scheduling modifications regulate the *timing, duration*, and *prioritization* of task execution.

Task timing. Task timing describes the time at which a task is initiated. Task timing does not always follow monotonically from task priority (Tulga & Sheridan, 1980). This is especially true if no tasks are considered truly urgent. Tasks usually occur in a sequence where certain preconditions (including the availability of information or resources or the completion of other tasks) must be satisfied before the task may commence and must be accomplished by some point in time, often dictated by the requirements of subsequent tasks. Task timing is a key challenge for humans working in complex, event-driven domains (Ho, Nikolic, Waters, & Sarter, 2004). Adams, Tenney, and Pew's (1994) review of the literature concludes that the cognitive management of multiple tasks requires the sequential scheduling of tasks because humans can truly work on only one task at a time; other tasks consequently must be queued. Task spin-up and spin-down place additional burdens on individuals in multitask environments (MacMillan, Deutsch, & Young, 1997). Moray, Dessouky, Kijowski, and Adapathya (1991) studied the effects of time constraints on timing tasks and found that humans do not time tasks optimally in the presence of time constraints because the time taken to determine optimal timing erodes the gains from adopting an optimal timing. This echoes the findings of Tulga and Sheridan (1980), who found that people do not plan ahead when very busy. Adaptations in task timing may schedule tasks with the goal of short-term optimization of resources to minimize slack time.

Task prioritization. Humans inherently place differential value on the tasks needed to accomplish their goals. A common prioritization scheme distinguishes task priority along two dimensions: urgency and importance (Covey, 2004). Urgency can be affected by factors such as time to respond and the certainty of the information, whereas importance can be determined by factors such as the level of threat or the potential impact of task failure on safety or mission objectives. Higher priority is usually given to those tasks that are both urgent and important. Changes to the task context can alter the urgency and importance of a task and accordingly alter its priority (Ho et al., 2004). Tasks that have a higher priority are often allowed to interrupt those that have a lower priority (Navon & Gopher, 1979). Adapting the priority of tasks may couple with other elements of a task scheduling scheme to change the timing of task execution, the order in which tasks are performed, or whether the task is even performed at all, given time and resource constraints.

Task duration. Tasks take time. The time each task requires is a combination of the nature of the task itself, the tools available, the human's experience performing the task, and other mediating contextual factors. Many tasks inherently have a finite window during which they must be accomplished. The time available to accomplish a task is subjective and highly variable, although some minimum time can usually be calculated if all needed information is available. Adaptations in time allocations may involve setting finite deadlines for the completion of a task or changing the time allotted once a task has been started.

Implications and considerations. A recognizable portion of task scheduling centers on interruption management, which involves reassessing the task timing and duration to account for a new prioritization scheme. Literature from a variety of domains has confirmed that humans are easily interrupted and that poor handling of interruptions can increase errors, increase frustration and stress, and reduce efficiency and decision quality (Chen & Vertegaal, 2004; Gillie & Broadbent, 1989; Iqbal, Adamczyk, Zheng, & Bailey, 2005; McFarlane & Latorella, 2002). McFarlane and Latorella (2002) conclude that a principled approach to improve interface design for interruption management is lacking, despite guidelines that recognize the operator's need for greater control of tasks, such as those proposed by Smith and Mosier (1986).

Adaptations to task scheduling aim to create automated systems that are less frustrating to operators by endowing them with the same courtesies exhibited by human colleagues, such as an understanding of task priorities and interruptibility (Bickmore, 2010; Miller & Parasuraman, 2007; Parasuraman & Miller, 2004). One of the challenges in designing effective interruption and task scheduling systems is the difficulty automated systems have predicting the interruptibility of a human operator. Having this ability would allow automated systems to take advantage of periods of high interruptibility to suggest the operator attend to a different program or switch tasks and use periods of low interruptibility to minimize disruptions. Although Adamczyk, Iqbal, and Bailey (2005) were able to create models to predict interruptibility with a 78% accuracy compared to self-reports based on physiological measures and a task model, work is needed in this area to improve modifications to task scheduling.

Examples. Alerting systems are good examples of automated systems that reprioritize tasks for humans and dictate the timing of certain tasks. Pritchett (2001) categorizes alerting systems into three main types: signal detectors, hazard detectors, and hazard resolvers. All three seek to interrupt normal operator activities and to draw the operator's attention. The roles that alerting systems play in the modification of task scheduling include task management aid, overlord, initiator of procedures, desired cue, trusted monitor, and attention director (Pritchett, 2001). Alerting systems in modern commercial aircraft actively modify their behavior depending on the phase of flight. For example, minor warnings are routinely suppressed during the takeoff and landing phases of flight so as not to interrupt the pilot while performing other critical flight tasks.

How: Modification of Interaction

The modification of interaction adapts the interaction between the human and the automation and seeks to answer questions such as how information is exchanged, where the interaction locus of control is, how often the operator will interact with the automation, and when this interaction will happen. Modifications of interaction are typically not recognized as a class of adaptations by function allocation-centric taxonomies. Modifications to interaction comprise adaptations to the *interface features, interaction style*, and the *amount of interaction* required.

Interface features. One of the simplest adaptations is the modification of the way the information is displayed to the operator. Here, we take a narrower view than Rothrock et al. (2002) and define modifications to interface features as only modifications to the information layout, ways to augment the information to direct attention (e.g., highlighting), and changes to the navigation of information (e.g., context-dependent menus).

Amount of interaction. The amount of interaction an individual has with automation is defined along two dimensions: how much interaction and when that interaction occurs. How much interaction spans a continuum from very little to continuous. The interaction between the human and the automation can come at any point in the task. Historically, the majority of interaction has come either at the beginning (activation) or the end (response selection) of a task. Alternatively, interaction with the automation can be interspersed throughout a task or purposefully interrupt a task.

Modality. The modality refers to the sensory channel (e.g., visual, auditory, haptic) in which information is exchanged (e.g., visual, auditory, haptic). Wickens's multiple resource theory contends that the dynamic allocation of incoming information to the most readily available attentional resource pool will avoid overtaxing the operator. Separate resources are defined by auditory versus visual processing and spatial (analog) versus verbal (linguistic) processing (Wickens & Hollands, 2000). Information can be presented in a different modality (e.g., visual to text, or text to speech).

Interaction style. The interaction style refers to the rules that govern the interaction employed by automation, how information is exchanged between agents, and the locus of authority and control of interaction (Billings, 1997). The first aspect of interaction style governs whether the information exchange between agents is either given or requested. The act of requesting information is not the same as the task of giving information and may result in differential performance of the joint human–machine system (Entin & Entin, 2001; Sperling, 2005). The second aspect of interaction style is the locus of control of the interaction, which can also be thought of as the automation's authority level. Authority level over the joint set of tasks has two extremes: full human authority or full automation authority. Two points on this continuum are better known as *management by delegation* (also called management by permission) and *management by exception* (Billings, 1996). Unlike modifications of function allocation, management by permission versus exception does not change who does the task (i.e., task assignment or function allocation) but rather who has final authority over the task. For instance, under management by delegation rules, one agent (traditionally the human) permits the automation to execute the task assigned to it; the locus of control remains with the agent who delegates tasks in real time.

Implications and considerations. Modifications to the display features are some of the most widely used adaptations and the considerations for use are well documented in the literature (see Rothrock et al., 2002, for a thorough review). However, adaptations to the other aspects of interaction, interaction amount and style, are not widely utilized by current adaptive systems. Changes to the amount of interaction may have been overlooked thus far or be seen to "fall out" from modifications of other dimensions. Changes to interaction style, on the other hand, have likely been actively avoided as changes to interaction style often violate the human factors tenant of consistent behavior. Operators work to construct mental models of the automation; thus, modifications to the interaction style have the potential to be disruptive. On the other hand, changes in the interaction style can be used to increase user engagement, a well-established practice in computer gaming (S. D. Whitlow, personal communication, November 8, 2011).

As authority and roles change between the human and the automation, another consideration is one of *automaton etiquette*, which facilitates smooth and effective interactions between people and automation (Hayes & Miller, 2010). A change in authority may carry implications for a change in the power and familiarity relationship between the human and the automation (Bickmore, 2010). In addition, how those relationships are expressed can be culturally dependent.

Examples. Many examples of adaptations to the interaction style can be found in educational fields in the form of tutoring or coaching systems. Several examples of adaptations in these areas include adapting to the learning styles of the student (Johnson et al., 2003), adapting the amount of interaction with the automated tutor as the expertise of the student changes (Dorneich & Jones, 2001), and adapting the changing relationship between coach and student (Bickmore, 2010).

Although adaptive systems that modify the locus of authority are rare, Barnes and Grossman (1985) identified three types of adaptations to control authority: emergency logic, executive logic, and automated display logic. Emergency logic modifies authority based on the automation's situation assessment. Emergency adaptations are taken without the consent of the human operator, thereby changing the interaction locus of control from the human to the automation. Executive adaptations, on the other hand, are taken with human approval and therefore retain the original interaction locus of control.

What: Modification of Content

Modification of content describes changes to the information content itself (as opposed to the modification of interaction category, which governs how the information is interacted with but where the information stays largely unchanged). Often, these types of changes are designed to provide a subset of all available information to support short-term goals and tasks. Changes to content can usually be described as changes to *quantity*, level of *abstraction*, or *quality*.

Quantity. The most straightforward adaptation to information content is to decide whether to present the information at all. Dynamic display decluttering adaptations are used to help operators focus on only the most important information for the task at hand.

Abstraction. Information can be aggregated or abstracted to focus on salient aspects and reduce the processing time for the human to interpret the presented information. Choosing the correct abstraction for the key functional relationships is often the aim of ecological interface design (Burns & Hajdukiewicz, 2004). As the functional relationships of interest change, so too should the information displayed to the operator.

Quality. Similarly, the information may be presented at lower or higher quality, depending on the current context. If information is of a time-critical nature, then the system may provide a lower quality version if the full information content would exceed the deadline for task completion. For example, video streams may be reduced in frame rate or even converted to a series of lower-quality images (Mohan, Smith, & Li, 1999).

Implications and considerations. Typically content adaptations are best suited for knowledge acquisition and analysis support tasks. Content adaptations are also suited to highlight key functional relationships that may need additional operator attention. When modifying the content, it is imperative to ensure that the user is provided with the information needed to successfully accomplish his or her work.

Examples. With the proliferation of mobile handheld devices, there has been significant work to modify content based on factors such as the client display capabilities, quality of service considerations, and network state (Shaha, Desai, & Parashar, 2001). Modifications typically include reducing the quality of images and video for lower-bandwidth transmission and aggregating the content into a simpler navigational scheme.

TAXONOMY OF TRIGGERS

By definition, adaptations are designed for specific situations as defined by context and as such, changes in context can trigger the system to adapt the automation. Triggers are based on several classes of information that can be sensed, observed, or modeled to create an understanding of context or "what is happening in the world" relevant to the adaptive system's decision making. An adaptive system needs triggers to identify when to engage an adaptation, how long an adaptation should persist, and when to disengage the adaptation. Previous

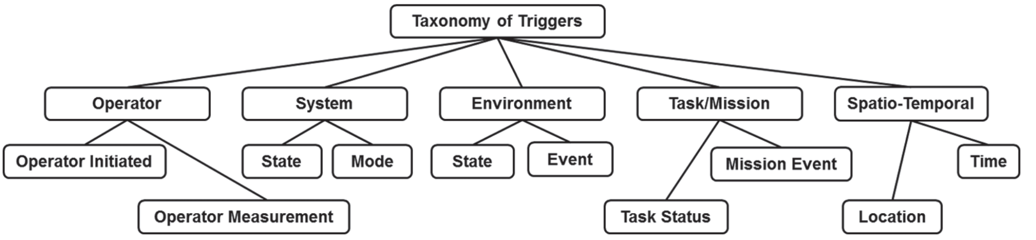

Figure 3. Taxonomy of Triggers for adaptive systems.

discussions of adaptation management have focused primarily on engagement triggers, with less discussion on duration and disengagement criterion, the importance of which has been highlighted with aviation accidents caused by unanticipated autopilot disengagement (Billings, 1996). Expanding on previous categorizations (Byrne & Parasuraman, 1996; Parasuraman, Mouloua, & Molloy, 1996; Rouse, 1988; Sheridan & Parasuraman, 2006), this taxonomy classifies adaptation triggers into five broad categories (see Figure 3): operator, system, environment, task, and spatiotemporal.

- *Operator-based triggers*. Adaptations can be triggered by the operator directly or by a system assessment of the operator state.
- *System-based triggers*. Current or predicted states of the system can be used to trigger adaptations. Different modes of system operations can also trigger adaptations.
- *Environment-based triggers*. States of the environment or events external to the operator and the system can be used to trigger adaptations.
- *Task- and mission-based triggers*. A mission is typically composed of a coherent set of goals and subgoals and accomplished by a set of tasks. Triggers can be based on task state or mission state.
- *Spatiotemporal triggers*. Both time and location can be used as adaptation triggers.

Often adaptive systems consider multiple triggers, used in conjunction, to identify when the operator needs additional (or reduced) automation support.

Operator-Based Triggers

The simplest and original method for adaptation management is human request, where human operators engage and disengage automation as needed or desired. Recently, researchers have been exploring direct means for measuring operator state using model-based or sensor-based information to enable automation to trigger adaptations.

Operator initiated. Who controls automation adaptations is a question that has received a fair amount of attention (Kaber, Riley, Tan, & Endsley, 2001; Miller & Parasuraman, 2007; Opperman, 1994; Parasuraman et al., 1996; Parasuraman, Mouloua, Molloy, & Hilburn, 1993; Prinzel et al., 2000; Scerbo, Freeman, & Mikulka, 2003), resulting in a distinction between those adaptations that are under human control (*adaptable*) and those that are not under human control (*adaptive*; Opperman, 1994). This article has included both types in its discussions without distinction since the Taxonomy of Adaptations is applicable to systems with both automation- and human-initiated triggers. Recent work on operator-initiated adaptive automation has included delegation methods where the human commander delegates tasks to automation as he or she might to a junior teammate (Arciszewski et al., 2009). Miller and Parasuraman (2007) maintain that delegation is inherently powerful because the supervisor (or human) can choose which tasks to delegate to automation, the method by which the task is to be accomplished, and how much monitoring, approving, and reviewing are required. Operator-initiated automation, however, is limited in scope as it requires direct operator input, time, and attention—which may be unavailable—to initiate any automated function.

Operator measurement. As the goal of adaptations is to improve joint human–machine system performance and human workload, measurement of these operator characteristics can be used to trigger adaptations. Other operator characteristics such as fatigue, visual load, and stress can also be measured. When direct measurement is not possible, estimations or models have been substituted (Scerbo et al., 2003).

The reduction (Huey & Wickens, 1993) and stabilizing (Miller & Parasuraman, 2007) of workload are major drivers for adaptive systems. If one could measure workload, a workload reducing automation could be designed to turn on when the user's workload is high to help shoulder some of the burden and to turn off when workload is low to avoid boredom. As physiological measures have become easier to obtain and research has shown them to be more reliable, it is now possible to use physiologically derived workload measures to drive adaptations (Bailey, Scerbo, Freeman, Mikulka, & Scott, 2006; Byrne & Parasuraman, 1996; Prinzel et al., 2000; Prinzel, Freeman, Scerbo, Mikulka, & Pope, 2003; Sharma, 2006; Wilson & Russell, 2007). Scerbo et al. (2003) assessed the positive and negative attributes of possible physiological measures and concluded that the most promising is EEG. Prinzel et al. (2000) showed how EEG could be used to provide a real-time assessment of workload using an engagement index as well as negative feedback to drive an adaptive system to successfully improve performance and decrease workload for simple tracking tasks. The success of EEG-based methods has led to an emphasis on the development of more robust EEG measurement devices and classification algorithms (Dorneich, Mathan, Ververs, & Whitlow, 2008).

Performance, the second widely measured operator parameter characteristic, is often seen as being roughly inversely proportional to workload (Wickens, 1992). Performance models aim to evaluate the human operator's present goals (including behavior) and his or her ability to perform tasks efficiently, maintain situation awareness, gauge information processing resources, and plan actions and goals (Gray, 2007; Rouse, 1988). They have been used in many adaptive systems (Benyon & Murray, 1993; Brusilovsky, 1996; Virvou, 1999; Wickens, 1992).

System-Based Triggers

Knowledge of the system is often used to support adaptation decisions. The system knowledge encapsulated in a system model may include its structure, modes, internal states, anticipated future states, and range of potential actions. Thus, system-based triggers include the *system state* and the *system mode*.

System state. System state is a description of the current configuration of the automation and is a relatively straightforward trigger (Parasuraman, Bahri, Deaton, Morrison, & Barnes, 1992). For example, in an automobile, system state includes position, speed, and acceleration. An adaptive system might use such information to increase steering sensitivity at high speeds, a feature found on many luxury cars.

System mode. System modes describe a grouping of several system configurations under one label where typically each mode corresponds to a set of unique system behaviors (Johnson, 1990). This has also been described as *functional modes* of a system (Degani, Shafto, & Kirlik, 1999). The human operator can monitor or select system modes, which are typically designed to help operators develop an understanding of automation behavior (e.g., automobile cruise control). Modes are often used in combination with other triggers to afford greater specificity.

Environment-Based Triggers

The environment can be modeled as a representation of the relevant facets of the world outside the immediate system and operator (Pritchett, Feigh, Kim, & Kannan, 2011). This model can be conceived of as a model of the work domain (Vicente, 1999). Often, environment models are knowledge-based and relatively static compared to the other models except in cases where the world is rapidly changing, which would generate situations that might require adaptations. Two categories of the environment-based triggers are *states* and *events*.

Environment state. States in the environment can be a description of the environmental parameters. Examples of environmental state triggers include changes in ambient light level (e.g., hand-held displays that change the display background color), temperature or humidity (e.g., programmable thermostats), or even wind speed (e.g., large wind turbines that feather their blades when wind speeds exceed safe velocities; Bossanyi, 2000).

Environment event. Adaptations can be triggered when external events occur in the environment. For example, an event might be an external threat event that triggers the automation of aircraft defensive measures following the detection of antiaircraft radar detection (Barnes & Grossman, 1985).

Task- and Mission-Based Triggers

A mission is typically composed of a set of tasks designed to achieve a set of goals, subject to constraints. Triggers can be based on *mission state* or *task state.*

Mission state. A mission is typically organized into phases or subgoals, each of which is subject to constraints such as the time to complete and pre- and postconditions. Many adaptive systems identify the need for adaptation by comparing expected actions to observed actions based on knowledge of the mission, plan, intent, or goals of the joint human–machine system. In addition, completion of mission phases can be used as the basis of engagement, persistence, and disengagement triggers. Mission-based triggers are specified at a higher level of abstraction than task-based triggers.

Task state. Adaptation management based on task state uses the initialization, completion, or partial completion of tasks (regardless of their impact on mission goals or objectives) to drive changes in automation. The challenge is how to identify which task the human is currently working on without requiring the human to continually inform the automation manually (Miller & Funk, 1997; Stiles, Bodenhorn, & Baker, 1998). Typically, this is done by monitoring the human's interaction with the automated system (Miller & Funk, 1997). For example, SEDAR is a CAD system to assist roofing designers that includes a critiquing agent that automatically adapts its comments to fit the current design task, as identified by monitoring the objects that the designer selects from SEDAR's menus (Fu, Hayes, & East, 2007). In another example, the Cockpit Information Manager Rotorcraft Pilot's Associate triggered cockpit configuration adaptations based on a model of active, planned, and completed tasks (Miller & Funk, 1997; Stiles et al., 1998).

Spatiotemporal Triggers

Time and location can be used to trigger adaptations.

Time based. The temporal criterion is a simple mechanism to manage the engagement and disengagement of automation. In the DFA literature, researchers have found the benefits (Hilburn, Molloy, Wong, & Parasuraman, 1993; Parasuraman, Hilburn, Mol, & Singh, 1991) and costs (Ballas, Heitmeyer, & Perez, 1992) of short-cycle versus long-cycle adaptive automation (where the system oscillated between manual and automatic control). Time triggers alone, however, are seen as having limited applicability (Byrne & Parasuraman, 1996; Prinzel et al., 2003).

Location based. The location of the automated system can be used as a trigger. The location can be absolute (e.g., GPS location of an aircraft) or relative (e.g., 10 miles from the top of descent in an aircraft flight path).

SUMMARY

Adaptive systems are those that automatically sense or track information from the environment, task, system, or user and adapt to help people to be more effective despite changing conditions. The drive to design effective adaptive systems is based on a desire to support the changing needs of people who perform highly complex work in dynamic environments. The introduction of adaptive systems has been a direct response to the need to provide multiple automation configurations that can be invoked based on the automation's assessment of specific contextual features—a move from point design to robust design.

Challenges and roadblocks to practical implementation of adaptive systems include the difficulty of implementing and controlling automation that can adapt itself and automatically and unobtrusively sense and interpret the user's cognitive state. Recent improvements, however, in neurophysiological and physiological sensing as well as explorations in how to interpret and use this information effectively have made adaptive systems more practicable and encouraged new applications and additional research.

The framework presented in this article provides a structure that can help researchers to organize a diverse range of literature on adaptive systems and help systems developers systematically consider the range of possible adaptations. What is known about the impact of adaptations and how best to trigger them has been uneven, with much more emphasis on modifications to function allocation and content.

The research and development community is still learning the implications and effects of adaptive systems, regardless of how well implemented. As such, many research questions remain open, including better understanding of the following:

- The nature of disruptions experienced by users caused by the specific types of adaptations or triggers remains unclear. What happens when the automation adapts in a way that surprises, confuses, irritates, or creates extra work for the user? Is it possible to avoid surprises and confusion, and if so, how?
- What are the inherent trade-offs caused by specific adaptations that system developers need to anticipate? For example, the ability to focus on a single safety-critical task may need to be achieved at the cost of situation awareness on the secondary tasks temporarily shouldered by the automation.
- What are the short-term and long-term implications of the changes in workload (both positive and negative) and performance caused by specific system adaptations and combinations of adaptations?
- What is the best way to leverage ever-improving approaches for measuring, interpreting, and using information on the user's cognitive state, by itself or in combination with other types of triggers given in the taxonomy?
- What are empirical methods and metrics to assess the reliability, accuracy, and robustness of adaptive systems? For example, how may one determine whether an adaptive system is adapting at the "right" time (Pritchett, Kim, Kannan, & Feigh, 2011) or in the right direction?
- Better understanding of the impact of task and context on cognition is needed, including a better theory of context (Hammond, 1996; Hammond, Hamm, Grassia, & Pearson, 1987).

Such future investigations will enable the community to understand why and when to use adaptive systems and how to create more practical, more accurate, more usable, less obtrusive, and more responsive user-centered systems to help us in our work and our lives.

ACKNOWLEDGMENTS

The authors would like to thank Chris Miller and William Rogers for extensive review of the article. In addition, the authors would like to acknowledge the efforts of colleagues who participated in related projects that informed much of the thinking of this article: Patricia May Ververs, Stephen D. Whitlow, Santosh Mathan, Jim Carciofini, Janet Creaser, Trent Reusser, and Jeff Rye.

KEY POINTS

- Unlike traditional human–automation systems that fix the roles of humans and automation at the design stage, adaptive systems aim to enhance joint human–automation performance by having the technological portion of the joint human–machine system invoke varying levels of automation support in real-time during task execution.
- Adaptive systems are defined here as having three key characteristics: a range of capabilities often configured into modes of operation, contextual awareness, and the authority to initiate changes to the system's functionality and operator interface.
- This article proposes a framework to categorize the two key elements of an adaptive system: the aspects of automated systems open to adaptation (Taxonomy of Adaptations) and the methods to trigger those adaptations (Taxonomy of Triggers).
- The Taxonomy of Adaptations categorized the wide range of methods to adapt automation into four areas: modification of function allocation, modification of interaction, modification of content management, and modification of task management.
- The Taxonomy of Triggers described the different methods by which adaptations could be triggered: operator, system, environment, task/mission, and spatiotemporal.

- The framework taxonomies provide a systematic way to organize research on specific adaptations or triggers.
- The framework presented here is a starting point for system designers by illuminating both the great opportunities afforded by the range of adaptations and the potential pitfalls that system designers must guard against.

REFERENCES

Abbott, T. S., & Rogers, W. H. (1993). Functional categories for human-centered flight deck design. In *Proceedings of 12th American Institute of Aeronautics and Astronautics—Institute of Electrical and Electronics Engineers Digital Avionics Systems Conference (DASC)* (pp. 66–74). Reston, VA: American Institute of Aeronautics and Astronautics.

Adamczyk, P. D., Iqbal, S. T., & Bailey, B. P. (2005). A method, system, and tools for intelligent interruption management. In *TAMODIA* (pp. 123–126). Gdansk, Poland: ACM Press.

Adams, M. J., Tenney, Y. J., & Pew, R. W. (1994). Situation awareness and the cognitive management of complex systems. *Human Factors, 37*, 85–104.

Alter, K. W., Goins, J. B., Hofer, E. F., Koehn, W. L., Miles, W. L., Mowry, R. S., & Pfaff, T. A. (1995). *High speed research flight deck design and integration flight deck concepts* (Tech. rep.) Seattle, WA: Boeing/McDonnell Douglas Industry Team.

Arciszewski, H., de Greef, T., & van Delft, J. (2009). Adaptive automation in a naval combat management system. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Human, 39*(6), 1188–1199.

Bailey, N. R., Scerbo, M. W., Freeman, F. G., Mikulka, P. J., & Scott, L. A. (2006). Comparison of a brain-based adaptive system and a manual adaptable system for invoking automation. *Human Factors, 48*, 693–709.

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

Ballas, J. A., Heitmeyer, C. L., & Perez, M. A. (1992). *Direct manipulation and intermittent automation in advanced* cockpits (Final rep. NRL/FR/5534–92-9375). Washington, DC: Naval Research Laboratory.

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

Benyon, D. R., & Murray, D. M. (1993). Applying user modeling to human-computer interaction design. *Artificial Intelligence Review, 6*, 43–69.

Bickmore, T. (2010). Etiquette in motivational agents: Engaging users and developing relationships. In C.C. Hayes & C.A. Miller (Eds.), *Human-computer etiquette: Understanding the impact of human culture and expectations on the use and effectiveness of computers and technology* (pp. 205–230). Abingdon, UK: Taylor & Francis.

Billings, C. E. (1996). *Human-centered aviation automation: Principles and guidelines* (Technical Memorandum 110381). Washington, DC: NASA.

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

Bossanyi, E. A. (2000). The design of closed loop controllers for wind turbines. *Wind Energy, 3*(3), 149–163.

Botargues, P. (2008). *Airbus AP/FD TCAS mode: A new step towards safety improvement*. Haren, Brussels: Eurocontrol.

Brusilovsky, P. (1996). Methods and techniques of adaptive hypermedia. *User Modeling and User-Adapted Interaction, 6*, 87–129.

Burns, C. M., & Hajdukiewicz, J. R. (2004). *Ecological interface design*. Boca Raton, FL: CRC Press.

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

Chen, D., & Vertegaal, R. (2004). Using mental load for managing interruptions in physiologically attentive user interfaces. In *Proceedings of CHI* (pp. 1513–1516). Gdansk, Poland: ACM.

Covey, S. R. (2004). *The seven habits of highly effective people* (9th ed.). New York, NY: Free Press.

Degani, A., Shafto, M., & Kirlik, A. (1999). Modes in human–machine systems: Review, classification, and application. *International Journal of Aviation Psychology, 9*(2), 125–138.

Dekker, S., & Woods, D. D. (2002). MABA-MABA or abracadabra? Progress on human-automation co-ordination. *Cognition Technology & Work, 4*, 240–244.

Dorneich, M. C., & Jones, P. M. (2001). The UIUC virtual spectrometer: A java-based collaborative learning environment. *Journal of Engineering Education, 90*(4), 721–728.

Dorneich, M. C., Mathan, S., Ververs, P. M., & Whitlow, S. D. (2008). Cognitive state estimation in mobile environments. In D. Schmorrow & K. Stanney (Eds.), *Augmented cognition: A practitioner's guide* (pp. 75–111). Santa Monica, CA: Human Factors and Ergonomics Society.

Dorneich, M. C., Mathan, S., Whitlow, S. D., Ververs, P. M., & Hayes, C.C. (2010). Etiquette considerations for adaptive systems that interrupt: Cost and benefits. In C. C. Hayes & C. A. Miller (Eds.), *Human-computer etiquette: Understanding the impact of human culture and expectations on the use and effectiveness of computers and technology* (pp. 289–322). Abingdon, UK: Taylor & Francis.

Dorneich, M., Ververs, P., Mathan, S., Whitlow, S. C. J., & Reusser, T. (2006). Neuro-physiologically-driven adaptive automation to improve decision making under stress. In *Proceedings of the Human Factors and Ergonomics Society Conference*. Santa Monica, CA: Human Factors and Ergonomics Society.

Duncan, J. (1986). Disorganisation of behaviour after frontal lobe damage. *Cognitive Neurophychology, 3*, 271–290.

Endsley, M. R. (1987). The application of human factors to the development of expert systems for advanced cockpits. In *Proceedings of the Human Factors Society 31st annual meeting* (pp. 1388–1392). Santa Monica, CA: Human Factors and Ergonomics Society.

Endsley, M. R., & Kaber, D. B. (1999). Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics, 42*(6), 462–492.

Entin, E. E., & Entin, E. B. (2001). Measures for evaluation of team process and performance in experiments and exercises. In *Proceedings of the 6th International Command and Control Research and Technology Symposium*. Washington, D.C.: Command and Control Research Program.

Fitts, M. P. (1951). *Human engineering for an effective air-navigation and traffic control system* (Tech. rep.). Washington, DC: National Research Council, Committee on Aviation Psychology.

Fu, M. C., Hayes, C. C., & East, E. W. (2007). SEDAR: Expert critiquing system for flat and low-slope roof design and review. *Journal of Computing in Civil Engineering, 2*(1), 60–68.

Gillie, T., & Broadbent, D. E. (1989). What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research, 50*, 243–250.

Gray, W. D. (2007). *Integrated models of cognitive systems*. Oxford, UK: Oxford University Press.

Hammond, K. R. (1996). *Human judgment and social policy*. New York, NY: Oxford University Press.

Hammond, K. R., Hamm, R. M., Grassia, J., & Pearson, T. (1987). Direct comparison of the efficacy of intuitive and analytical cognition in expert judgment. *IEEE Transactions on Systems Man and Cybernetics, 17*, 753–770.

Hancock, P. A., Chignell, M. H., & Loewenthal, A. (1985). An adaptive human–machine system. In *IEEE 1985 Proceedings of the International Conference on Cybernetics and Society* (pp. 627–630). New York, NY: Institute of Electrical and Electronics Engineers.

Hayes, C. C., & Miller, C. A. (2010). Should computers be polite? In C. C. Hayes & C. A. Miller (Eds.), *Human-computer etiquette: Cultural expectations and the design implications they place on computers and technology* (pp. 1–14). Abingdon, UK: Taylor & Francis.

Hilburn, B., Molloy, R., Wong, D., & Parasuraman, R. (1993). Operator versus computer control of adaptive automation. In *7th International Symposium on Aviation Psychology*.

Ho, C.-Y., Nikolic, M. I., Waters, M. J., & Sarter, N. B. (2004). Not now! Supporting interruption management by indicating the modality and urgency of pending tasks. *Human Factors, 46*, 399–409.

Huey, B. M., & Wickens, C. D. (1993). *Workload transition: Implications for individual and team performance*. Washington, DC: National Academy Press.

Iqbal, S., Adamczyk, P. D., Zheng, X. S., & Bailey, B. P. (2005). Towards an index of opportunity: Understanding changes in mental workload during task execution. In *Proceedings of CHI* (pp. 311–320). Gdansk, Poland: ACM.

Johnson, J. (1990). Modes in non-computer devices. *International Journal of Man-Machine Studies, 32*, 423–438.

Johnson, W., Shaw, E., Marshall, A., & LaBore, C. (2003). Evolution of user interaction: The case of agent Adele. In *Proceedings of the 8th International Conference on Intelligent User Interfaces* (pp. 93–100). Gdansk, Poland: ACM.

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. R. (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 Ergonomic Science, 5*(2), 113–153.

Kaber, D. B., Riley, J. M., Tan, K.-W., & Endsley, M. R. (2001). On the design of adaptive automation for complex systems. *International Journal of Cognitive Ergonomics, 5*(1), 37–57.

Kaber, D. B., Wright, M. C., Prinzel, L. J., & Clamann, M. P. (2005). Adaptive automation of human–machine system information-processing functions. *Human Factors, 47*, 730–741.

Kaber, D. B., Wright, M. C., & Sheik-Nainar, M. A. (2006). Investigation of multi-modal interface features for adaptive automation of a human–robot system. *International Journal of Human-Computer Studies, 64*, 527–540.

Lee, S. M., Kim, S. Y., & Feigh, K. M. (2009). Structural framework for performance-based assessment of ATM systems. In *Aviation Technology, Information, and Operations Conference*. Hilton Head, SC: American Institute of Aeronautics and Astronautics.

Lintern, G. (2012). Work-focused analysis and design. *Cognition Technology & Work, 14*(1), 71–81.

MacMillan, J., Deutsch, S. E., & Young, M. J. (1997). A comparison of alternatives for automated decision support in a multitask environment. In *Proceedings of the Human Factors and Ergonomics Society 41st annual meeting* (pp. 190–195). Santa Monica, CA: Human Factors and Ergonomics Society.

Mazaeva, N., Dorneich, M. S. W., Mathan, S., & Ververs, P. (2005). Characterization of changes in electrophysiological activity in an operational environment. In *Proceedings of the Human Factors and Ergonomics Society conference* (pp. 1177–1181). Santa Monica, CA: Human Factors and Ergonomics Society.

McFarlane, D. C., & Latorella, K. A. (2002). The scope and importance of human interruption in human-computer interaction design. *Human-Computer Interaction, 17*(1), 1–61.

McGuire, J. C., Zich, J. A., Goins, R. T., Dwyer, J. B. E. J. P., Cody, W. J., & Rouse, W. B. (1991). *An exploration of function analysis and function allocation in the commercial flight domain* (Tech. rep., NASA Contractor rep. 4374). Washington, DC: NASA.

Meister, D. (1985). *Behavioral analysis and measurement methods*. New York, NY: John Wiley.

Miller, C. A., & Funk, H. B. (1997). Knowledge requirements for information management: A rotorcraft pilot's associate example. In M. Mouloua & J. M. Koonce (Eds.), *Human-automation interaction: Research and practice* (pp. 186–192). Mahwah, NJ: Lawrence Erlbaum.

Miller, C. A., & Hannen, M. D. (1999). The rotorcraft pilot's associate: Design and evaluation of an intelligent user interface tool for cockpit information management. *Knowledge-Based Systems, 12*(8), 443–456.

Miller, C. A., & Parasuraman, R. (2007). Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Human Factors, 49*, 57–75.

Mohan, R., Smith, J. R., & Li, C.-S. (1999). Adapting multimedia internet content for universal access. *IEEE Transactions on Multimedia, 1*(1), 104–114.

Moray, N., Dessouky, M. I., Kijowski, B. A., & Adapathya, R. (1991). Strategic behavior, workload, and performance in task scheduling. *Human Factors, 33*, 607–629.

Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review, 86*(3), 214–255.

Norman, D. A. (1986). *User centered design: New perspectives on human–computer interaction: Cognitive engineering*. Hillsdale, NJ: Lawrence Erlbaum.

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

Palmer, E. (1995). Ooops, "it didn't arm": A case study of two automation surprises. In *Proceedings of the Eighth International Symposium on Aviation Psychology* (pp. 227–232).

Parasuraman, R., Bahri, T., Deaton, J. E., Morrison, J. G., & Barnes, M. (1992). *Theory and design of adaptive automation in aviation systems* (Progress rep. NAWCADWAR-9023-60). Warminster, PA: Naval Air Warfare Center.

Parasuraman, R., & Bowers, J. C. (1987). *Psychophysiology of the electronic workplace: Attention and vigilance in human computer interaction*. New York: John Wiley.

Parasuraman, R., Hilburn, B., Mol, R., & Singh, I. (1991). *Adaptive automation and human performance: III. Effects of practice on the benefits and costs of automation shifts* (Tech. rep.). Warminster, PA: Naval Air Warfare Center.

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

Parasuraman, R., Mouloua, M., & Molloy, R. (1996). Effects of adaptive task allocation on monitoring of automated systems. *Human Factors, 38*, 665–579.

Parasuraman, R., Mouloua, M., Molloy, R., & Hilburn, B. (1993). Adaptive function allocation reduces performance cost of static automation. In *7th International Symposium on Aviation Psychology* (pp. 37–42).

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

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (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*(3), 286–298.

Price, H. (1985). The allocation of functions in systems. *Human Factors, 27*, 33–45.

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.

Prinzel, L. J., Freeman, F. G., Scerbo, M. W., Mikulka, P. J., & Pope, A. T. (2003). Effects of a psychophysiological system for adaptive automation on performance, workload, and the event-related potential p300 component. *Human Factors, 45*, 601–613.

Pritchett, A. R. (2001). Reviewing the role of cockpit alerting systems. *Human Factors and Aerospace Safety, 1*, 5–38.

Pritchett, A. R., Feigh, K. M., Kim, S. Y., & Kannan, S. (2011). *Work models that compute to support the design of multi-agent socio-technical systems*. Manuscript submitted for publication.

Pritchett, A., Kim, S. Y., Kannan, S. K., & Feigh, K. (2011). Simulating situated work. In *IEEE Conference on Cognitive Methods in Situation Awareness and Support* (pp. 66–73). New York, NY: Institute of Electrical and Electronics Engineers.

Rothrock, L., Koubek, R., Fuchs, F., Haas, M., & Salvendy, G. (2002). Review and reappraisal of adaptive interfaces: Toward biologically inspired paradigms. *Theoretical Issues in Ergonomic Science, 3*(1), 47–84.

Rouse, W. (1988). Adaptive interfaces for human/computer control. *Human Factors, 30*, 431–488.

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: Lawrence Erlbaum.

Scerbo, M. W., Freeman, F. G., & Mikulka, P. J. (2003). A brain-based system for adaptive automation. *Theoretical Issues in Ergonomic Science, 4*(1–2), 200–219.

Scerbo, M. W., Freeman, F. G., Mikulka, P. J., Parasuraman, R., DiNocero, F., & Prinzell, L. J. I. (2001). *The efficacy of psychophysiological measures for implementing adaptive technology* (Tech. Rep. NASA/TP-2001-211018). Hampton, VA: NASA Langley Research Center.

Shaha, N., Desai, A., & Parashar, M. (2001). Multimedia content adaptation for QOS management over heterogeneous networks. In *International Conference on Internet Computing* (pp. 642–648). Las Vegas, NV: CSREA Press.

Sharma, S. (2006). Linear temporal characteristics of heart interbeat interval as an index of the pilot's perceived risk. *Ergonomics, 49*(9), 874–884.

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

Sheridan, T. B., & Parasuraman, R. (2006). Human-automation interaction. *Reviews of Human Factors and Ergonomics, 1*, 89–129.

Sheridan, T. B., & Verplank, W. L. (1978). *Human and computer control of undersea teleoperators* (Tech. Rep. 780815025). Arlington, VA: Office of Naval Research.

Smith, S. L., & Mosier, J. N. (1986). *Guidelines for designing user interface software* (Tech. Rep. ESD–TR-86-278). Bedford, MA: MITRE.

Sperling, B. K. (2005). *Information distribution in complex systems to improve team performance* (Unpublished doctoral dissertation). Georgia Institute of Technology, Atlanta.

Stiles, P., Bodenhorn, C., & Baker, B. (1998). Decision aiding on rotorcraft pilot's associate. In *Proceedings of the Annual Forum of American Helicopter Society* (Vol. 54, pp. 1212–1224). Alexandria, VA: American Helicopter Society.

Tenney, Y. J., Rogers, W. H., & Pew, R. W. (1995). *Pilot opinions on high level flight deck automation issues: Toward the development of a design philosophy* (Contractor Rep. 4669). Hampton, VA: NASA Langley Research Center.

Tulga, M. K., & Sheridan, T. B. (1980). Dynamic decisions and work load in multitask supervisory control. *IEEE Transactions on Systems, Man and Cybernetics, 10*(5), 217–232.

Vicente, K. J. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Lawrence Erlbaum.

Virvou, M. (1999). Automatic reasoning and help about human errors in using an operating system. *Interacting with Computers, 11*, 545–573.

Wickens, C. D. (1992). *Engineering psychology and human performance* (2nd ed.). New York, NY: HarperCollins.

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

Wickens, C. D., Li, H., Santamaria, A., Sebok, A., & Sarter, N. B. (2010). Stages and levels of automation: An integrated meta-analysis. *Human Factors and Ergonomics Society Annual Meeting Proceedings, 54*, 389–393.

Wilson, G. F., & Russell, C. A. (2007). Performance enhancement in an uninhabited air vehicle task using psychophysiologically determined adaptive aiding. *Human Factors, 49*, 1005–1018.

Woods, D. (1996). *Decomposing automation: Apparent simplicity, real complexity*. Mahwah, NJ: Lawrence Erlbaum.

Karen M. Feigh is an assistant professor at the Georgia Institute of Technology's School of Aerospace Engineering. She earned a PhD in industrial and systems engineering from the Georgia Institute of Technology in 2008.

Michael C. Dorneich is a principal research scientist in the Human Centered Systems Group at Honeywell Laboratories. He earned his PhD in industrial engineering in the Human Factors Program at the University of Illinois at Urbana-Champaign in 1999.

Caroline C. Hayes is a professor at the University of Minnesota's Department of Mechanical Engineering. She earned a PhD in robotics from Carnegie Mellon University in 1990.

Date received: November 8, 2010
Date accepted: March 4, 2012