Mitigation Strategy Design for Optimal Augmented Cognition Systems

Susan Harkness Regli, Patrice Tremoulet, Helen Hastie, and Kathleen Stibler

Lockheed Martin Advanced Technology Laboratories
Cherry Hill, NJ 08002
sregli@atl.lmco.com

Abstract

The development of robust, accurate, real-time cognitive state gauges, based upon operators' physiological data, is an essential pre-requisite to realizing performance improvements through augmented cognition technology. However, gauges that can identify undesirable cognitive states are only useful if corrective action can be taken when the gauges detect problematic states. Thus, in designing augmented cognition systems, it is equally essential to develop mitigation strategies that help operators avoid and/or move out of cognitive states that can cause degradations in operator performance.

This report addresses three aspects of mitigation strategy design that need to be considered when deciding what systems should do in response to changes in operator cognitive state: 1) interaction of mitigation strategies with general human-computer interaction principles; 2) task-domain dependency; and 3) correlation between mitigation strategies and specific combinations of gauge values. The treatment of each issue is grounded in specific examples from the design and implementation of multiple augmented cognition systems.

1 INTRODUCTION

Research in augmented cognition has made progress in developing reliable cognitive state gauges that can give insight into when an operator is entering undesirable cognitive states for successful job performance. Early on, the assumption was that if the gauge output was available in real-time and was accurate, changes to the operator's work environment could be made to re-establish an optimum cognitive state for performance. These environment changes, in the form of human-computer interface adjustments, are called "mitigation strategies."

In our augmented cognition research, we developed a technological approach and system architecture, named Performance Augmentation through Cognitive Enhancements (PACE), that can use gauge data to assess operator cognitive state and initiate mitigation strategies to present new information at an optimal pace and in appropriate modalities to optimize overall performance. The appropriate mitigations are triggered when the operator is in or is approaching a cognitive state that leaves the operator vulnerable to poor performance. The PACE system turns off the mitigations when the gauges indicate the user is no longer in a vulnerable state.

In the course of developing this technology, we realized that identifying the correct mitigation strategies to use in response to a particular reading from gauge outputs was not as simple as it sounded. Adjusting non-optimal cognitive states can be at least as difficult as recognizing them. We have identified three primary issues that need to be considered in the design of mitigation strategies. In this paper we will discuss each of the issues and present examples of how we have addressed these challenges that are inherent in mitigation strategy design.

2 MITIGATION STRATEGY DESIGN CHALLENGES

2.1 Mitigation Strategies and Human-Computer Interaction (HCI) Principles

Interaction between mitigation strategies and general HCI principles arises because of the following dilemma. On the one hand, mitigation strategies should be designed to be turned on or off when necessary to maintain acceptable levels of human resource loading and, thus, reduce negative stress on the operator. However, if an interface design succeeds in reducing operator stress, why shouldn't that interface strategy be used consistently, rather than only turning it on when the operator is overloaded? To capitalize on the ability of augmented cognition technology to detect and react to fluctuations in operator needs, mitigation strategies that have benefits during periods of high stress but potentially unacceptable costs during normal operations must be identified.

The remainder of this section describes mitigation strategies that were designed specifically to exhibit the right cost/benefit balance. The key is in recognizing how cognitive challenges vary between normal conditions and periods of high stress. The examples in this section draw primarily on work in which cognitive state gauges were being developed to assess working memory usage, in particular as broken down into verbal working memory and spatial working memory taxation levels.

2.1.1 Intelligent Sequencing

The intelligent sequencing mitigation strategy ensures that secondary task activities are scheduled to support maximum operator performance by timing the presentation of verbal and spatial secondary tasks when gauges indicate that the operator is not overloaded verbally or spatially, respectively.

For example, as operators perform primary tasks that are a mix of both spatial and verbal tasks (as they are in most operational environments) and secondary tasks arrive, PACE (Morizio, Thomas & Tremoulet, 2005) classifies them as mostly verbal or mostly spatial. Meanwhile, spatial and verbal working memory gauges indicate when the subject's verbal and/or spatial memory stores are taxed by the primary task. Based on secondary task classifications and the gauge readings, PACE schedules the secondary tasks so that they interrupt primary tasks on an optimum schedule. For example, if gauges indicated that an operator's spatial working memory store is taxed, only verbal secondary tasks will be presented until memory taxation has stabilized.

The cost of intelligent sequencing of secondary tasks and alerts is that the operator does not have access to total situational information at all times (i.e., low priority information and tasks are kept on hold until the high-stress situation resolves). This removes some of the operator control over his own task environment and creates a risk that information will be missed. The benefit during high-stress periods is that more information and tasks should be able to be processed over time and the operator's cognitive resources will be maximally optimized.

2.1.2 Pacing

The Pacing mitigation strategy involves directing tasks according to an operator's workload and arousal levels, as determined by arousal and cognitive workload gauges. Research has indicated that the timing of an interruption relative to a user's current task load can affect the user's ability to cope with the interruption (Czerwinski et al., 2000; McFarlane, 2002; Monk, Boehm-Davis & Trafton, 2002). The Pacing mitigation builds on the Intelligent Sequencing mitigation by: 1) adding the identification of appropriate "cognitive breaks" to the total information determining when tasks should be presented and 2) allowing primary tasks to be decomposed so tasks and subtasks can be optimally scheduled at a higher granularity than just primary versus secondary tasks.

For example, if arousal and workload gauges indicate that the operator's task load is above maximum threshold, pending tasks are queued for delivery rather than presented immediately. The pending tasks are delivered when the gauges indicate a break in cognitive activity. If too many secondary tasks are competing for delivery, the primary task will be decomposed into subtasks so that cognitive breaks (during which secondary tasks can be delivered) occur more frequently. Overall throughput of memory-taxing tasks is optimized.

Pacing mitigation strategy exhibits all of the same costs and benefits of the Intelligent Sequencing mitigation (Section 2.1.1). An additional risk of the Pacing strategy is that the task decomposition may not match the operator's understanding of the task structure well enough to support natural processing and behavior. An additional benefit of the pacing mitigation is that the operator can multitask at a lower level of task granularity to make maximum use of time on task.

2.1.3 Modality-based Task Switching

This strategy involves developing alternate display strategies to invoke specific sensory modalities. Based on output from sensory gauges (e.g., verbal and spatial working memory), as well as an understanding of current system state (i.e., which verbal and spatial tasks are currently being performed and their relative loading, or task intelligence), display strategies that invoke modalities with spare capacity and/or that are best suited for the information to be communicated are employed. For example, while spatial information (e.g., location of a threat) is generally best presented via visual imagery (e.g., target on radar screen), an operator's capacity for using spatial information might be overly taxed, increasing the possibility that the information about the threat might be missed. In this case, the threat location might be more effectively presented using another modality: it could be presented as sound localization (e.g., auditory signal at a given location) or as tactile cues (e.g., vibration of a sensor in a tactile vest). The use of alternate modalities for presenting the information about the threat will increase the likelihood that the operator will process and respond to important situational information.

The costs of Modality-based Task Switching are primarily related to the number of input channels an operator has to attend to simultaneously to receive all relevant situational information. During low-stress periods, an operator might want all information to be delivered using a single modality so that attention and monitoring efforts can be focused. During high-stress periods, however, the Modality-based Task Switching mitigation can take advantage of all modalities to which a human has capacity to attend.

2.2 Generic versus Application-Specific Mitigation Strategies

Another challenging aspect of mitigation design is determining how much information a system should have about operator tasks to effectively use mitigation strategies. From the standpoint of system architecture and implementation, generic strategies have strong benefits in their ability to be reused and validated across multiple domains. Accordingly, initial efforts to develop mitigation strategies for augmented cognition systems focused upon defining generic mitigation strategies. However, recent research suggests that human performance may increase more when deployment of mitigation strategies is controlled not only by cognitive state assessments from physiological data but also by information about the operator's current task context. This section contrasts a generic mitigation strategy, chunking, with an application-specific mitigation, displaying a custom-designed summary panel. Both strategies are intended to mitigate against performance declines caused by high cognitive workload, but while the first was inspired by classic research in cognitive psychology, the second was the product of task analyses and design reviews conducted to facilitate integrating the PACE architecture with a new application—a prototype user interface for a future release of the Tactical Tomahawk Weapons Control System (TTWCS).

2.2.1 Chunking

While chunking is generally thought of as grouping of information based solely upon semantic content (Miller, 1956), such groupings may also be encouraged via modality differences. Nelson and Bolia (2003) demonstrated that spatially grouping audio displays can enhance auditory-cue identification by approximately 50 percent, as well as speed recognition response times. In this work they placed call signs at different auditory spatial locations (i.e., each call sign sender was assigned a specific spatial location).

Chunking may also be done according to criticality. For instance, in high stress, high workload situations, the method of delivering new information to an operator could be modified to help focus the operator's attention upon the most critical set(s) of information. That is, knowing a user is currently overloaded (based on cognitive workload assessment), an augmented cognition system could stop using ordinary methods (traditional output modality and information format) to convey new information to the operator. Alerts, status changes and other data that are outside a predefined critical context can be queued, e.g., in a small message window. (Critical and non-critical messages in the queue should be visually distinguished and the operator should be able to sort the messages based on whatever criteria are relevant for the application.) Meanwhile, the system will immediately inform the operator of any new high-priority alerts related to the critical task context through alternative interruption strategies that facilitate chunking, such as sound localization or vibrations from a tactile vest.

2.2.2 Displaying a Summary Panel

Task analyses and design reviews of the v6.0 TTWCS user interface prototype uncovered the following potential problem. The interface relies on multiple action tabs (Strike Overview, Route Planning, Prepare Missile and Post-launch) that change color when a mission element requires an action, such as accepting a revised plan or sending a report. The operator has to click on the highlighted tab to see what mission needs attention and what action needs to be taken. During periods of high cognitive workload (e.g., when multiple missions are in progress), each of the primary action tabs on the interface are likely to display an indicator that there is an “Action Required” on some aspect of that action tab. The result is that multiple tabs will indicate a need for attention at the same time, and the operator will have to decide which to attend to first. This decision making process could add cognitive work to an already overloaded operator.

The tab metaphor is a perfectly appropriate interface element in this system for periods of low workload, because the display allows the operator to focus on one activity at a time. It may be inappropriate, however, during periods of high workload. During those periods, the additional activity of prioritizing which actions to attend to in which order may result in the operator clicking back and forth to determine the most critical actions that need to be taken.

In this particular application, the mitigation design is a visual decision-making aid that can be displayed in response to high workload gauge readings to mitigate against performance declines. More specifically, when the cognitive state gauge indicates that operator cognitive workload is high, a summary panel will appear over the bottom left portion of the screen, which normally would be available for the expansion of the tables under the action tabs.

This summary panel will give a quick “at-a-glance” summary of information that needs attention in all of the action tabs, with each action associated with a short piece of information, such as the mission number, to allow the operator to assess the relative priority of each action (see Figure 1).

Strike Overview		Route Planning		Missile		Post Launch
51011	Send Exc	51003	Send SCO	51001	Make Ready	
51012	Send Exc			51002	Make Ready	
				51003	Make Ready	

Figure 1. Summary panel details showing mission number and action required.

The operator will have the ability to move and minimize the summary panel. The summary panel will disappear when the operator’s cognitive workload drops below the “high” threshold or when the operator minimizes the window.

The summary panel display mitigation and the chunking mitigation are both true user interface design mitigations; that is, both involve manipulating the way information is presented to users based on their cognitive state and the current task context. However, they differ in terms of application specificity. Chunking may be used as a strategy to mitigate against high workload with any of a wide variety of applications. In contrast, the summary panel display strategy is directly customized for the v6.0 TTWCS user interface prototype. Of course, this mitigation could be made more generic. However, even if it were rewritten as a summary window that lists the most relevant details of each of the views in a multi-view display, it could still only be applied to a relatively small set of applications: namely, those with interfaces designed such that operators may, in some high workload situations, need to decide which of multiple alternate views to attend to first to perform the tasks at hand. On the other hand, this mitigation has great potential to significantly improve operator performance for these types of applications. The costs associated with applying the summary panel display mitigation during normal operations include increasing screen clutter and distracting the user; tabs are a good design technique for low and moderate workload because they help to declutter, save screen real estate, and focus operator’s attention. Displaying the summary panel during high workload benefits operators by assisting them in deciding what view to attend to first.

2.3 Cognitive State to Mitigation Strategy Mapping

Because mitigation strategies are triggered by the output of gauge values, a critical challenge in mitigation strategy design is the correlation of the correct mitigation strategy with the correct gauge output. The challenge is complicated by the need to use multiple gauges to account for the complex nature of human cognition. In one example, using the output from only one gauge (workload) resulted in negative consequences for experiment participants. The workload gauge was triggering the pacing mitigation strategy, which meant that if workload was low, more tasks were to be delivered. During the test some participants became overwhelmed with a high workload and effectively "gave up," but when the workload gauge went down in response to their giving up on the tasks, even more tasks were delivered. Hence just as participants stopped trying to perform tasks, they were given more to do, resulting in even more work left undone (Tremoulet & Regli, in process).

Multiple gauges can help avoid this type of problem; in the example above, the combination of a workload gauge with a gauge for arousal or engagement would have indicated a reason to avoid giving the overloaded participants more work (i.e., their workload is low, but they are also not engaged in the tasks). To effectively use multiple gauges, however, two design challenges must be met. First, readings from multiple gauge outputs must be mapped to practical descriptions of cognitive state. Second, once the mapping to cognitive state is complete, appropriate mitigation strategies must be assigned to each cognitive state based on what is needed to establish, maintain, or regain optimal cognitive state. If gauges report values for both working memory usage and arousal, for example, mitigation strategy designers need to decide what it means when spatial working memory is taxed, but arousal is low: the operator may be trying to complete many working memory tasks while lethargic. In this case the mitigation needs not only to reduce the memory taxation but also to get the operator's attention back on task. Ensuring that the mitigation matches the real cognitive state depicted by all gauges is critical to avoid applying inappropriate mitigation strategy and inadvertently causing more harm than good.

The remainder of this section provides an example of a gauge-to-mitigation matrix developed to determine viable mitigation strategies to respond to the outputs of five different cognitive state gauges: Arousal, Workload, Total Working Memory (WM), Verbal WM, and Spatial WM. The two-step process was extremely complex given the number of gauges, but the outcome yielded only a few areas of real uncertainty where contradictory readings from gauges made it difficult to predict an appropriate mitigation technique.

In the first step (Table 1), the output of the arousal and workload gauges was combined to determine a cognitive state description that related to alertness and engagement in the tasks at hand. Similarly, output of the Working Memory gauges was combined to determine a cognitive state description that related to memory usage and taxation. Once the initial cognitive state descriptions were determined, the arousal/workload states and the working memory states were combined to determine an overall description of cognitive state based on the outputs of all five gauges. The outcome of this process could be used by an augmented cognition system to automatically label cognitive states to prepare for triggering a mitigation strategy.

In the second step (Table 2), the arousal/workload cognitive state descriptions were associated with mitigation strategies that could be used to re-establish optimum cognitive state. (Background calculations of desired outcomes were conducted to assist in the development of effective mitigations.) Similarly, the cognitive state descriptions were associated with appropriate mitigation strategies to maximally utilize memory resources. Once the initial mappings were complete, both sets of mitigations were compared to determine the appropriate course of action for each set of five gauge readings.

Three possible outcomes occurred during this comparison and are reflected in the final output: 1) the working memory mitigations and the arousal/workload mitigations were consistent and could be simultaneously employed; 2) a more severe mitigation is needed to override a less severe mitigation (e.g., if one of the cognitive state descriptions indicated a problem and the other did not, the problem indicator prevailed); or 3) the mitigation strategies conflicted. In the last case, we assumed that more work would need to be done to determine if the inconsistent readings ever occurred simultaneously, and if so, under what conditions.

3 CONCLUSION

The lessons learned from addressing the mitigation strategy design challenges described in this paper have highlighted the fact that research in augmented cognition has progressed to a point where we can begin to truly attend to the integration of cognitive state gauge output in the operational environments of operators. Our research community has begun to answer the question: "What can we know about dynamically changing cognitive states?" But now we face an equally challenging secondary question: "What can we do about it?"

The examples presented here represent work to date in recognizing and beginning to answer this question. The research and experiences highlight the need to attend to three primary issues when developing appropriate responses

Table 1. Gauge Output Mapped to Cognitive State Descriptions

GAUGES			Arousal	OK	TOO LOW	TOO LOW	TOO LOW	BORDERLINE	TOO HIGH	BORDERLINE	OK	OK	TOO HIGH	TOO HIGH	BORDERLINE
			Workload	OK	OK	BORDERLINE	TOO HIGH	OK	OK	TOO HIGH	TOO HIGH	BORDERLINE	BORDERLINE	TOO HIGH	BORDERLINE
Total Working Memory	**Verbal Working Memory**	**Spatial Working Memory**	**Cognitive States**	**Okay**	**Lethargic**	**Lethargic or Given Up**	**Given Up (Because Critically Overloaded)**	**Distracted / Emotionally Taxed**	**Emotionally Overloaded**	**(Assume) Overloaded**	**Distracted or Near Given Up**	**Assume OK**	**(Emotionally) Overloaded**	**Overloaded**	**Possible Overload**
OK	OK	BORDER-LINE	**WM untaxed**	WM untaxed	lethargic	assume lethargic	given up	distracted / emotionally taxed	emotionally overloaded	assume near overload	distracted	assume OK	(emotionally) overloaded	assume overloaded	assume near overload
BORDER-LINE	OK	OK	**assume distracted**	assume distracted	assume lethargic or may be about to give up	assume lethargic	overloaded and given up	distracted	emotionally overloaded	assume near overload	near given up	assume OK	(emotionally) overloaded	overloaded	assume OK
TOO HIGH	OK	OK	**assume distracted**	assume distracted	assume distracted and lethargic	assume given up	overloaded and given up	distracted	emotionally overloaded	overloaded	given up	distracted	(emotionally) overloaded	overloaded	assume overloaded
OK	BORDER-LINE	BORDER-LINE	**assume WM taxed**	assume WM taxed	assume lethargic	assume lethargic	given up	assume near overload	emotionally overloaded	assume near overload	near given up	assume WM taxed	(emotionally) overloaded	overloaded	near WM overload
BORDER-LINE	BORDER-LINE	BORDER-LINE	**WM taxed**	WM taxed	assume lethargic; possible overload	assume lethargic	given up	near overload	overloaded	overloaded	near overload; near given up	WM taxed	overloaded	overloaded	near WM overload
BORDER-LINE	TOO HIGH	TOO HIGH	**assume WM overloaded**	assume WM overloaded	assume given up	assume given up	given up	overloaded	overloaded	overloaded	near overload; near given up	assume WM overloaded	overloaded	overloaded	assume WM overloaded
TOO HIGH	BORDER-LINE	BORDER-LINE	**WM overloaded**	WM overloaded	assume given up	given up	given up	overloaded	overloaded	overloaded	overloaded; near given up	assume overloaded	overloaded	overloaded	overloaded
TOO HIGH	TOO HIGH	TOO HIGH	**WM overloaded**	WM overloaded	assume given up	given up	given up	overloaded	overloaded	overloaded	overloaded; near given up	overloaded	overloaded	overloaded	overloaded

Table 2. Cognitive State Descriptions Mapped to Mitigation Strategies

Cognitive States	**Arousal / Workload**	**Okay**	**Lethargic**	**Possible Overload**	**(Assume) Overloaded**	**Possible Overload**	**Distracted**	**Given Up / Critically Overloaded**	**Lethargic or Given Up**	**Possible Overload**
Working Memory	**Mitigations**	**None**	**Get Attention Mitigation**	**Give Option to Defer**	**Defer and Give Option to Delegate Work**	**Give Option to Delegate Work**	**Chunking or Modality Switch**	**Delegate Work**	**Give Option to Defer or Delegate Work or Get Attention Mitigation**	**Get Attention Mitigation or Delegate Work (Consult Other Gauges)**
okay	**none**	none	get attention mitigation	give option to defer	defer and give option to delegate work	give option to delegate work	chunking or modality switch	delegate work	get attention mitigation	get attention mitigation
WM untaxed	**accept work**	all accept work	get attention mitigation	give option to defer	defer and give option to delegate work	give option to delegate work	chunking or modality switch	delegate work	get attention mitigation	get attention mitigation
assume WM overloaded	**give option to defer**	all give option to defer	give option to defer	give option to defer	defer and give option to delegate work	defer and give option to delegate	chunking and give option to defer	delegate work	give option to defer and give option to delegate	delegate work
assume WM overloaded (distracted)	**give option to delegate work**	all give option to delegate work	give option to delegate work	give option to delegate work	defer and give option to delegate work	give option to delegate work	chunking and give option to delegate	delegate work	delegate work	delegate work
assume WM overloaded	**chunking or modality switch**	all chunking or modality switch		chunking/modality switch and give option to defer	chunking/modality switch and defer	chunking/modality switch and give option to delegate	chunking or modality switch	delegate work	delegate work	delegate work
WM overloaded	**defer**	all defer		defer work	defer and give option to delegate work	defer and give option to delegate work	chunking/ modality switch and defer	delegate work	defer and give option to delegate work	delegate work
WM overloaded	**delegate work**	all delegate work		delegate work	delegate work	delegate work	chunking/ modality switch and delegate	delegate work	delegate work	delegate work

Consistent mapping beween Working Memory and Arousal/Workload mitigations

A more severe mitigation overrides a less severe mitigationA more severe mitigation overrides a less severe mitigation

Mitigation strategies conflict with one another to undesirable cognitive states in an operational setting:

1) How can mitigation strategy design identify and account for the differences in cognitive challenges during undesirable cognitive states (arousal levels that are too high or too low, for example) versus during optimal states and normal operations?
2) What level of system knowledge of operator tasks is optimal to balance between power of implementation versus the benefits of domain specificity?
3) What is the appropriate mapping, at very low levels of granularity, from cognitive gauge output to cognitive state, and then from cognitive state to actual mitigation response.

We trust that as research in augmented cognition progresses, many more issues such as these will be identified. The ongoing challenge is to take mitigation strategy design seriously as the critical "last mile" in successfully augmenting human cognition in operational environments.

4 ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of Dylan Schmorrow, Amy Kruse, Colby Raley, and Lee Kollmorgen. The work reported here was funded by the Defense Advanced Research Projects Agency (DARPA), under contract NBCHC030032 to Lockheed Martin Advanced Technology Laboratories.

5 REFERENCES

Czerwinski, M., Cutrell, E., & Horvitz, E. (2000). Instant messaging: Effects of relevance and time. In S. Turner & P. Turner (Eds.), *People and computers XIV: Proceedings of HCI 2000* (pp. 71-76). Sunderland, England: British Computer Society.

McFarlane, D. C. (2002). Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction*, 17(3), 312-338.

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

Monk, C., Boehm-Davis, D., & Trafton, J. G. (2002). The attentional costs of interrupting task performance at various stages. *Proceedings of the 2002 Human Factors and Ergonomics Society Annual Meeting*. Santa Monica, CA: HFES.

Morizio, N., Thomas, M., & Tremoulet, P. (2005). Performance augmentation through cognitive enhancement (PACE). *Proceedings of the HCI International Conference*. Las Vegas, NV: HCI International.

Nelson, T. & Bolia, R. (2003). Evaluating the effectiveness of spatial audio displays in a simulated airborne command and control task. *Proceedings of the human Factors and Ergonomics Society 47th Annual Meeting* (pp. 202-206). Denver, CO: Human Factors and Ergonomics Society.

Tremoulet, P. & Regli, S. (2006). Augmented cognition for command and control operators. *Journal of Cognitive Engineering and Decision Making (JCEDM).* (in review).