Towards a Closed-Loop Training System: Using a Physiological-Based Diagnosis of the Trainee's State to Drive Feedback Delivery Choices

Amy Bolton[1], Gwendolyn Campbell[2], and Dylan Schmorrow[3]

[1] Strategic Analysis Inc., 3601 Wilson Blvd., Suite 500, Arlington, VA 22201
[2] NAVAIR Orlando Training Systems Division, 12350 Research Pkwy., Orlando, FL 32826
[3] Office of Naval Research, 875 Randolph St., Suite 1425, Arlington, VA 22203
abolton@sainc.com, Gwendolyn.campbell@navy.mil,
Schmord@onr.navy.mil

Abstract. Designers of a closed loop scenario based training systems must have specifications to drive the decisions of whether or not performance feedback is appropriate in response to student behavior, the most effective content of that feedback, and the optimal time and method of delivery. In this paper, we propose that physiological measures, when interpreted in conjunction with information about the learning objective, task environment and student performance, could provide the data necessary to inform effective, automated decision processes. In addition, we present an overview of both the relevant literature in this area and some ongoing work that is explicitly evaluating these hypotheses.

Keywords: Simulation based training, physiological measures, feedback, training interventions.

1 Introduction

Optimal delivery of instruction is both critical and challenging in dynamic, scenario-based training (SBT) computer simulations such as those used by the military. Tasks that human instructors must perform during these sorts of simulated training exercises can impose a heavy burden on them. Partially due to advances in the state-of-the-art in training technology and partially due to the military's desire to reduce the number of personnel required, it may be possible to support functions that overburdened instructors perform by automating much of the SBT process in a closed-loop computer simulation. Unfortunately though, after more than 50 years of literature documenting research conducted in the area of training interventions, few empirically-supported guidelines have emerged to direct the choice and implementation of effective, automated training interventions including the types of measures that should be captured to inform the interventions. Designers of a closed loop SBT system must have specifications to drive the decisions of whether or not feedback is appropriate in response to student behavior, the most effective content of that feedback, and the optimal time and method of delivery. In this paper, we propose that physiological

D.D. Schmorrow, L.M. Reeves (Eds.): Augmented Cognition, HCII 2007, LNAI 4565, pp. 409–414, 2007.

measures, when interpreted in conjunction with information about the learning objective, task environment and student performance, could provide the data necessary to inform effective, automated decision processes. In addition, we will present an overview of both the relevant literature in this area and some ongoing work that is explicitly evaluating these hypotheses.

2 Feedback: Why and What?

One of the biggest challenges in training is making the inferential leap from student behaviors and performances that can be observed to the underlying knowledge, skills and abilities (KSAs) that are presumed to generate those outputs. We want to be able to make valid inferences because the "why" to giving feedback is, typically, to correct deficiencies in those underlying KSAs. Unfortunately, as linguists have been pointing out for decades, performance is an imperfect indicator of underlying capability. A correct action does not necessarily imply that the student has the required KSAs. Similarly, an error does not necessarily imply that the student does not. A set of behaviors, especially behaviors exhibited during a training exercise, is often a conglomeration of guesses, intentional actions and maybe even a few slips (Norman, 1981; Reason, 1990). If there was some way to sort these all out and recognize which actions were guesses, which were intentional, and which were accidental, instructors and training systems would be in a much better position to deliver the most effective feedback.

Recent advances in neurophysiological assessment suggest that EEG data may provide one source of insight into the specific nature of a trainee's individual actions and help distinguish deliberate and learned responses from exploratory "guesses" and unintentional slips. For example, previous fMRI and EEG research suggests that certain incorrect actions are followed by a negative deflection with a mediofrontal distribution that peaks at about 50-150 ms after the action (Luu & Pederson, 2004; Luu, Tucker & Makeig, 2004). This response is called the "error related negativity" or ERN response. Research has shown that this response does not seem to follow deliberately committed errors (Dehaene, Posner & Tucker, 1994; Stemmer, Witzke & Schönle, 2001). Instead, it appears as if the ERN follows actions that are incongruent with cognitive expectancies (Luu & Pederson, 2004; Luu, Tucker & Makeig, 2004). In other words, this response appears to be an "Oops, that's not what I meant to do" signal, making it a candidate measure for discriminating between the two types of errors, "slips" and "mistakes." Similarly, ongoing research is converging on the finding that there are several indicators that systematically index the learning process, such as the lateral inferior frontal negativity (LIFN), medial-frontal negativity (MFN) and P300 (Luu, Tucker & Stripling, manuscript in preparation). These responses are candidates for distinguishing guesses from learned responses.

Knowing why a student requires feedback (or some form of follow-on training and/or practice) is a good first step in determining what that feedback should be. While any instructional system could be improved upon, a reasonable first approximation of design specifications might go something like this: guesses should be remediated with feedback designed to fill in missing KSAs; true mistakes should be remediated with feedback designed to correct faulty and/or incomplete KSAs; and slips should be remediated with more practice, designed to establish automated procedures. Note that neurophysiological data alone, however, are not sufficient to generate instructional content. Rather, they help us sort behaviors into “like” categories, in preparation for further analyses; and these further analyses are based on an understanding of the learning objectives, the task environment and objective measures of student performance.

This relationship can be illustrated by describing an ongoing effort to conduct an assessment of the potential value added of incorporating neurophysiological measures into a SBT system. Subjects will be trained to categorize objects within the context of a warfare-based simulation. For any given object, the information available to support the categorization may be incomplete and/or ambiguous. In some cases, the environment will eventually and naturally provide information about the accuracy of a prior categorization decision, but this will not always be the case. (These features are not uncommon in real world learning environments.) Mathematical techniques, such as regression or fuzzy logic, can be applied to a set of data containing a subject’s categorization decisions for a number of objects and the information available to the subject about each object when the decision was made.

Campbell and colleagues (2006) have shown that it is possible to interpret the resulting mathematically-expressed decision rules, critique them by comparing them to rules derived from expert performance data, and generate qualitative feedback that, at least in some cases, leads to improved performance. They also showed, unfortunately, that error variance in the data set often posed a significant challenge to the modeling process. There were, of course, brute force approaches available to reduce this effect. First, they threw out a significant chunk of the data collected during the early learning trials, in order to reduce the “noise” generated by guesses and allow the “signal” to be detected by the algorithm. Even after taking that step, however, they had to have their participants complete many scenarios, and make many decisions, in order to avoid accidentally finding spurious patterns in the performance data. The fact that the some of the environmental cues were either ambiguous or correlated (or both) resulted in a situation in which a few data points representing slips had the potential to mislead the modeling process, unless they were outweighed by a very large number of intentional decisions.

If EEG indicators, such as the P300 and ERN, could have been used to separate intentional decisions from slips and guesses, it is possible that more accurate models of participants’ decision making rules could have been built earlier in the training process, allowing tailored and adaptive feedback to be delivered to each participant more quickly. And it is just this hypothesis that is being tested under the current effort. Two, randomly assigned groups will complete the scenario-based training. One group will have their decision making data modeled following the techniques used in Campbell, et. al. (2006). The other group will have their decision making data filtered using EEG indicators, before applying the modeling algorithms. The efficiency of each technique will be evaluated based on the amount of performance data that must be collected before a reliable model can be built and model-based feedback can be delivered. The effectiveness of each technique will be evaluated based on the steepness of the learning curve across the course of the training. This is the first effort that we are aware of to investigate the capability of neurological data to help answer the feedback questions of “why?” and “what?”.

3 Feedback: When and How?

There are a couple of efforts underway investigating whether or not neurophysiological measures should and could inform the feedback questions of "when?" and "how?" The basis for this research includes a recently completed investigation by Bolton (2006) and one that is underway (Van Buskirk, in preparation). In these initial studies, the research focused on determining, if one had an accurate diagnosis of trainee cognitive state (perhaps supported by neural and physiological data), when and how would one intervene with feedback? Therefore, in both of these investigations, training scenarios were explicitly designed a priori to elicit particular cognitive states from the trainee during a simulation based training exercise. The elicited cognitive states targeted by the scenario design were consistent with the hypotheses of each of the investigations. If this research proved promising, the next step would be to integrate measurement technologies and algorithms that could provide neural and physiological markers of cognitive state in real-time in order to drive the feedback timing and feedback methods.

In a doctoral dissertation, Bolton (2006) sought to provide empirical guidance for the optimal timing of feedback delivery (i.e., immediate vs. delayed) in a dynamic, scenario based training (SBT) computer simulation. As a basis for the investigation, the argument was presented that a general theory of feedback timing was likely not possible and that 50 years of literature on the topic was overly simplified. To deal with this presumed oversimplification, Bolton (2006) provided an integration of theoretical perspectives to provide a more detailed and sensitive feedback timing hypothesis. First, Bolton presented the premise underlying temporal contiguity, that contiguous feedback with performance encourages accelerated learning of cue-strategy associations (Guthrie, 1935). Second, the theory of transfer appropriate processing was presented that states that the more similar the conditions are between the training setting and the testing setting, the more positive transfer of training will be produced (Morris, Bransford, & Franks, 1977). Third, Bolton likened dynamic, SBT simulations to a complex, multi-tasking environments which suggested that care should be taken in understanding the cognitive demands placed on the trainee, consistent with the tenets of Cognitive Load Theory (CLT; Sweller, 1993; Sweller, Van Merrienboer, & Paas, 1998). Finally, the conclusion was drawn that the benefits of temporal contiguity of feedback should be balanced against the potential costs of changing the requirements of the task or disrupting task performance during training with the presentation of immediate feedback.

In other words, in a dynamic, SBT simulation system, temporal contiguity of feedback to decisions and actions is important. However, to prevent the feedback from serving as a task interruption, the cognitive load of the scenario could be used to indicate the appropriate timing of the feedback delivery. If the cognitive load was low enough, immediate feedback could and should be provided. If the cognitive load was too high, the delivery of feedback should be delayed until the end of the scenario. The hypotheses developed for the investigation specifically addressed expected variations in performance and instructional efficiency consistent with the above postulation.

In order to test the hypotheses, training scenarios were designed a priori to represent either low cognitive load or high cognitive load and feedback was delivered either immediately during a scenario or delayed until the end of a scenario. This created 10 experimental conditions (a partial factorial design) and 120 volunteers were randomly assigned to one of those conditions. After familiarization on the experimental testbed, participants completed a total of seven, 10-minute scenarios, which were divided across two training phases. During each training phase participants received either immediate or delayed feedback and performed either high or low cognitive load scenarios. Four subtask measures were recorded during test scenarios as well as subjective reports of mental demand, temporal demand and frustration.

A series of planned comparisons were conducted to investigate the training effectiveness of differing scenario cognitive loads (low vs. high), timing of feedback delivery (immediate vs. delayed), and sequencing the timing of feedback delivery and the cognitive load of the scenario. In fact, the data did not support the hypotheses as tested. However, through post hoc exploratory data analyses, some light was shed on the possibility that even the more detailed feedback timing statement, that immediate feedback should be provided during low cognitive load training scenarios and feedback should be delayed when the cognitive load of the training scenario is high, might have been over simplistic. The data seemed to suggest that the nature of the task, the phase of the training and the cognitive state of the trainee should be taken into consideration when considering effective feedback timing strategies. Specifically, different tasks imposed different cognitive demands and those demands changed with additional training as expertise was developed. Additionally, the a priori specification of scenario cognitive load is difficult to define. At any point in time during a single training or testing scenario, cognitive load demands likely vary. Different sub-tasks likely place different demands on the brain.

Based on this premise, Van Buskirk (in preparation) designed an investigation to determine if the delivery of real-time feedback could be optimized during scenario-based training so as not to place additional demands on areas of the brain that were already approaching overload – the “how” for the chosen “when”. This is consistent with evidence of multiple resource pools and memory stores (Baddeley & Logie, 1999; Wickens & Holland, 2000). Considering the constructs of spatial working memory and verbal working memory, the question is if real-time feedback can be delivered in a way so that is does not interfere with resource pools involved in task execution and learning. For example, if one is performing a highly taxing visuo-spatial training task, could feedback provided in the auditory channel be more effective than feedback provided to the visual channel or no feedback provided at all. This investigation is still underway therefore the results are still “to be determined.”

4 Summary

The literature and investigations discussed in this paper point to the assertion that to realize the vision of a closed-loop training system, the “why,” “what,” “when,” and “how” of feedback delivery, would require a sensitive measure of varying cognitive demands and states in real-time during training. Neurophysiological measures of cognitive state assessment could potentially provide the solution needed to realize this vision. With a more sensitive measure of what the trainee has learned and hasn’t learned and with a more sensitive measure of cognitive load as it changes over time, more sensitive training interventions could be researched and developed. The result of incorporating neurophysiological measures into a training system could be optimized, closed-loop training system that produces maximal training effectiveness and efficiency. At this time, the authors are aware of at least one ongoing effort at Clemson University that has the goal of including physiological measures in a replication of the experiment conducted by Bolton to gain a better understanding of how such measures may benefit training.

References

1. Bolton, A.E.: Immediate versus delayed feedback in simulation based training: Matching feedback delivery timing to the cognitive demands of the training exercise. Unpublished doctoral dissertation. University of Central Florida, Orlando, FL
2. Campbell, G.E., Buff, W.L., Bolton, A.E.: Viewing training through a fuzzy lens. In: Kirlik, A. (ed.) Adaptation in Human-Technology Interaction: Methods, Models and Measures, pp. 149–162. Oxford University Press, Oxford (2006)
3. Dehaene, S., Posner, M., Tucker, D.: Localization of a neural system for error detection and compensation. Psychological Science 5(5), 303–305 (1994)
4. Guthrie, E.R.: The Psychology of Learning. Harper, New York (1935)
5. Luu, P., Tucker, D., Makeig, S.: Frontal midline theta and the error-related negativity: Neurophysiological mechanisms of action regulation. Clinical Neurophysiology 115, 1821–1835 (2004)
6. Luu, P., Tucker, D.M., Stripling, R.: Neural mechanisms underlying the learning of actions in context (manuscript in preparation)
7. Luu, P., Pederson, S.: The anterior cingulate cortex: Regulating actions in context. In: Posner, M.I. (ed.) Cognitive neuroscience of attention, pp. 232–244. Guilford Publication, New York (2004)
8. Morris, C.D., Bransford, J.D., Franks, J.J.: Levels of processing versus transfer appropriate processing. Journal of Verbal Learning and Verbal Behavior 16, 519–533 (1977)
9. Norman, D.A.: Categorization of action slips. Psychological Review 88, 1–15 (1981)
10. Reason, J.: Human error. Cambridge University Press, Cambridge (1990)
11. Stemmer, B., Witzke, W., Schönle, P.W.: Losing the error related negativity (ERN): an indicator for willed action. Neuroscience Letters 308(1), 60–62 (2001)
12. Sweller, J.: Some cognitive processes and their consequences for the organisation and presentation of information. Australian Journal of Psychology 45, 1–8 (1993)
13. Sweller, J., Van Merrienboer, J.J.G., Paas, F.G.W.C.: Cognitive architecture and instructional design. Educational Psychology Review 10, 251–296 (1998)
14. Van Buskirk, W. L.: The use of feedback in simulation based training: Investigating the relationship between feedback timing, content, and modality under high cognitive workload. Unpublished doctoral dissertation. University of Central Florida, Orlando, FL (in preparation)