

THE ROLE OF TEMPORAL SEQUENCE LEARNING IN GUIDING VISUAL ATTENTION ALLOCATION

Sarah M. Miller & Wai-Tat Fu
Human Factors Division and The Beckman Institute
University of Illinois at Urbana-Champaign
Urbana, Illinois

Models of visual attention allocation suggest that monitoring is driven primarily by proximal cues like bandwidth and value. However, these cues might not always be predictive of the meaningful events an operator is asked to monitor. The aim of the current study is to extend visual sampling models by studying whether sampling can be influenced by more distal cues, like detecting patterns in the monitored signal, when proximal cues, like bandwidth, are not predictive of the meaningful events the operator is asked to monitor. Ten participants completed a task based on Senders' (1964) experiment where operators were asked to monitor a series of four gauges to detect when the gauges traveled into the alarm region. The performance results suggest that participants could successfully adapt to the temporal sequence. However, participants did not show explicit awareness of the sequence, indicating that this type of learning could, in some cases, be implicit. Implications for display design and training are discussed.

INTRODUCTION

What factors guide visual attention allocation in a monitoring task? Several models have been developed to predict how operators scan or monitor displays based on multiple cognitive and environmental factors including signal characteristics, display layouts, and mental models. Understanding how these factors guide visual sampling is important for numerous applications including display design and operator training.

Senders (1964) developed one of the first and most influential models of visual attention allocation, suggesting that optimal scanning should be driven by the rate of information change, or bandwidth of the display. Operators sample areas of interest in order to reduce the uncertainty about the state of the system and to reconstruct the signal. Since bandwidth creates the uncertainty, the model predicts signals with higher bandwidth will be sampled more often than those with lower bandwidth. Senders found that much of the variance in operators' visual scanning was indeed explained by the bandwidth of the instrument.

Although bandwidth has been found to be a good predictor of visual attention allocation, especially for well-learned monitoring tasks, it has always been assumed that bandwidth is perfectly correlated with the meaningful events the operator is asked to monitor. For example, if operators were asked to monitor a series of gauges to detect when one of them alarmed, the area with the highest bandwidth also alarmed the most frequently. Therefore, it was unclear whether operators were adapting to the bandwidth itself, or to the alarm rate, since both were perfectly correlated. Recent models

decoupled bandwidth and alarm rate to determine if operators sample based solely on bandwidth (i.e. Senders' model), or if bandwidth was merely a proximal cue used for predicting alarms (Miller, Kirlik, Kosorukoff, and Byrne, 2004). It was found that the correlation between the bandwidth and the meaningful events does indeed influence monitoring. This suggests that it is not only bandwidth that drives sampling, but sampling is also mediated by how useful the cue is in predicting alarm rate.

The purpose of the current study was to further extend these sampling models by studying whether visual sampling can also be influenced by more distal cues, like detecting patterns in the monitored signal, when proximal cues, like bandwidth, are not predictive of the meaningful events the operator is asked to monitor (e.g. alarms). We created scenarios that have a repeating temporal sequence of alarms and tested whether people were able to adapt to this regularity in the environment to improve performance when other more proximal cues were not available.

Temporal sequence learning

There has been a long history of research on sequence learning (e.g., Cohen, Ivery, Keele, 1990; Marcus, Karatekin, & Markiewicz, 2006; Nissen & Bullemer, 1987). In a typical experiment, subjects were presented with a sequence of lights and asked to respond by pressing a corresponding sequence of keys. A pattern of lights recurs regularly within this sequence. It was found that subjects press the keys for repeated sequences faster than a randomly presented sequence. The faster response suggests that subjects have learned the patterns.

Recently, Marcus et al. (2006) showed that the anticipatory eye movements that shift attention to likely stimulus locations prior to stimulus onset could explain part of the faster reaction time.

We extend previous studies on sequence learning to a visual monitoring task. There are two major differences between the traditional studies and the current task. First, in visual monitoring tasks, all stimuli were visible at all times, and people were required to constantly shift their attention overtly to different locations to detect changes. Second, the durations between successive events (e.g., alarms) are variable, so that one needs to learn only *where* but also *when* an event will happen.

In reality, there are often regularities in the relationship among when the meaningful events happen in different areas of interest on a visual display. For example, the pattern of the gauges traveling into the alarm region often follows some regular pattern. This pattern could serve as a useful tool in visual monitoring. Indeed, recent research shows that eye movements in skilled behavior are highly regular, and the regularities emerge as a consequence of adapting to the inherent constraints and regularities in the task environment (Reichle & Laurent, 2006). However, it is unclear whether operators are able to detect patterns in multiple distal events and to what extent these patterns can be recognized in a dynamic task.

METHODS

Participants

Ten students from the University of Illinois at Urbana-Champaign participated in the study. Each participant completed a 90-minute study session and was paid \$8 per hour. In addition to hourly pay, the best performer received a \$10 bonus.

Display Design and Configuration

A sample experimental display is depicted in Figure 1. The display consists of four round gauges, one at each corner of the monitor. This display took up the entire screen space on a 17" flat panel monitor. The diameter of each gauge was 5.25 cm. The horizontal separation between the center of the gauges was 17.5 cm, and the vertical separation was 13 cm. Participants were positioned so that their eyes were 35 cm from the monitor, resulting in a visual angle between the gauges of 28.1 degrees horizontally and 21 degrees vertically. This visual angle was chosen so that eye movements were required in order to sample the gauges.

For each gauge, the dark gray area was the alarm or out of bounds region and the light gray was the safe range. When the trial begins, the pointers (shown at 12

o'clock in Figure 1) for each of the four gauges began to move back and forth. The participant's goal was to monitor each of the gauges to detect when a pointer traveled into the alarm region (i.e. dark gray). A correct detection was counted when a button on the keyboard corresponding to the out of bounds gauge was pressed within one second of the pointer passing the threshold.

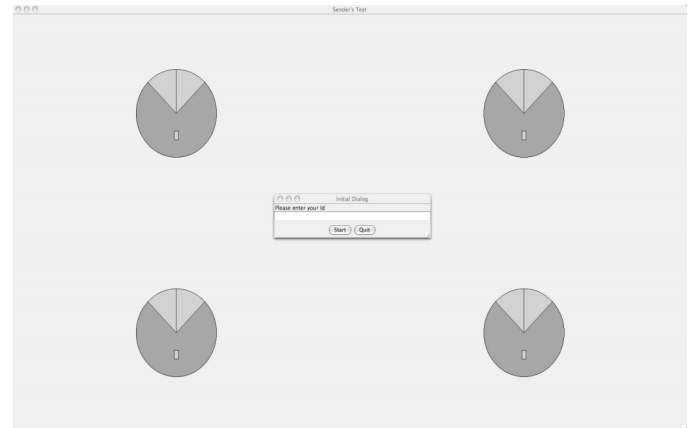


Figure 1. Experimental task at trial onset

In order to minimize the potential that proximal cues shown to be predictive in past experiments (e.g. Miller et al, 2004; Wickens et al, 2002) could be used to successfully guide sampling, the maximum bandwidth, or rate of change, was the same for each gauge (1.87 Hz). The total bandwidth summed across all four gauges was 7.5 Hz. This bandwidth was chosen based on the Nyquist sampling theorem so that operators would be required to continuously monitor the display but not be overloaded by the requirements (see Senders, 1964).

Additionally, cue validity, the relationship between the bandwidth and alarm rate was set to zero. This means the bandwidth was not predictive of alarm rate. Alarm rate, the number of times a pointer traveled into the alarm region per unit time, was also the same as in Senders (1964). Summed across the four gauges, alarm frequency was set to approximately 0.25 alarms per second. However, due to the reasons described in the following paragraphs, both cue validity and alarm frequency were not exactly 0 and 0.25 respectively.

The signal for each gauge was driven by the sum of three sine functions with a specific frequency, phase, and amplitude. In order to find a set of four such signals that satisfied all the experimental constraints, the frequency, phase, and amplitude parameters for each of the three sine functions comprising the signals were found through optimization. A genetic algorithm was developed specifically for this study to find a set of signals that satisfied all the constraints (Holland, 1992). This algorithm was designed to search a computational space in order to find a set of sine functions that

minimized the Euclidean distance from a specified target point (as defined by the constraints).

There were two signal sets generated (A and B), each having approximately the same alarm rate, BW, and cue validity. Two sets were developed as one method to test for sequence learning. This method will be described in greater detail in the experimental design section. Since finding an exact solution requires extensive computations, we chose to use a constraint-satisfaction heuristic search method (genetic algorithm) to find the approximate solutions. As a result, while the BW was the same for both signal sets, the alarm rate for both signal sets was only approximately 0.25 alarms/second and the cue validity (correlation between the BW and alarm rate) was only approximately 0. The actual values for the alarm rates for signal set A and B were 0.27 and 0.26 respectively. The actual cue validities were -0.02 and -0.05 respectively. These small differences were not expected to play a significant role in observed sampling differences. The only significant difference between the two signal sets was the temporal pattern of alarms.

Since the cues found to guide visual attention allocation in past experiments were set to be either equal or uninformative (e.g. bandwidth, value, alarm rate, cue validity), we were able to examine the potential of temporal patterns as a cue for detecting alarms. The primary difference between the two signal sets was in the pattern of the order in which the gauges alarmed. An example of an alarm pattern would be the following: upper-right, lower-left, lower-right, upper-right, upper-left, upper-left, lower-right, lower-left. The pattern takes approximately 30 seconds to complete, so each sequence was seen around 10 times per 5-minute trial.

Experimental Design and Procedure

The ten participants were randomly divided into two groups of five. Each group completed ten 5-minute trials. The first eight trials and trial ten were on a single signal set. The ninth trial was on the second signal set (thus serves as the “random” trial). The first group saw primarily signal set A, and only one trial with signal set B. In order to ensure that one signal sets was not inherently more difficult than the other, the second group was exposed primarily to signal set B, with one trial on set A.

Following the last trial, participants were shown sixteen, 30-second video clips. Following each clip, participants were asked to predict the gauge(s) that were to travel into the alarm region within the next second using either a four-option multiple-choice questionnaire (primarily group A) or a free response questionnaire (primarily group B). Between zero and three gauges could be out of bounds within the next second. Twelve

clips were from the trial seen nine times (trials 1 - 8, trial 10), while four clips were taken from the trial seen only once (trial 9). Finally, participants were asked if they noticed anything different about any of the trials and if they detected any patterns in the gauges.

RESULTS

Performance Measures

Performance measures were calculated as a score based on the manual (i.e. keyboard) responses to the alarms. Participants received one point for correctly detecting when a gauge traveled to the alarm region. A correct detection was counted if the alarm was detected within one second of traveling into the alarm region. If the alarm was not detected within a second, a miss was counted as minus one point. If the participants detected an alarm when there wasn't one (i.e. false alarm), a point was deducted from the total score. Total score (%) was calculated as the number of hits minus the number of misses and false alarms divided by the total number of alarms. Since both false alarms and misses were counted as negative points, it was possible for participants to receive a negative score.

Figure 2 shows the average total score (%) across all the trials for both signal groups. As shown in Figure 2, scores improved significantly throughout the study for both groups ($F(9,72)=24.654$, $p<0.01$). Additionally, scores were significantly worse in T9 than the previous trial for both groups ($t(9) = 5.062$, $p<0.01$; $t(9) = 4.252$, $p < 0.05$). There were no performance differences between groups ($F(1,8)=0.745$, $p=0.413$), suggesting that the significantly worse performance in trial 9 was due to changing the temporal sequence of alarms rather than an inherent difference in difficulty between the conditions.




Figure 2. Total score (%) performance results across 10 trials for groups A and B. The solid line is group A, while the dashed line is group B. Trial 9 (second-to-last) is the alternate sequence.

Questionnaire Results

For the post-test multiple-choice questionnaire based on the short video clip, twelve trials were taken from the primary signal set and four trials were taken from trial 9 (changed pattern). For each trial, a video clip from the task was run for approximately 30 seconds. The video clip was stopped so that all gauges were approximately equal distances from the alarm region. This means that the proximity to the alarm region could not be used as a cue to predict the gauge(s) that were about to alarm. Participants were asked to either select the next gauge that was going to alarm (group A) or select the gauge(s) that were going to alarm within the next second (group B). Participants in group B were told they could select between 0 and 3 gauges, since 3 was the maximum number of gauges that could alarm within the next second.

The purpose of the questionnaire and video clips was to determine if participants could explicitly recognize the patterns in the alarms. If indeed patterns could be recognized, it would be expected that participants would do better than chance at indicating which gauge(s) were to alarm in the next second. For group A, chance performance was 25%, since participants were asked to predict which one of four gauges that were about to alarm. For group B, chance performance was 6.7%, since there were fifteen potential combinations of potential answers. In addition to hits, misses and false alarms were also counted.

For group A (multiple-choice), the average percent correct for the twelve trials taken from the primary signal set was 69% (standard error = 11.4%). The average percent correct for the four trials taken from the secondary (trial 9) signal set was 25% (standard error = 16.0%). For group B (free-response), the average hit rate was 43%, while there were an average of 6.4 false alarms (16 questions). Performance on the questionnaire was significantly better than chance ($t(4) = 2.17, p < 0.05$)

We also compared performance in the task with performance on the posttest questionnaire. Performance in the task was measured as the average total score in the final two pre-transition trials (trials 7 and 8). Performance in the posttest questionnaire was measured as the total percent correct (group A), or hit rate (group B). The correlation between task and questionnaire performance was -0.04. The potential interpretations and implications of these results will be discussed in the next section.

DISCUSSION

The purpose of the current study was to determine if operators recognize and utilize temporal patterns in the gauge alarms when more proximal cues,

like signal bandwidth, were not informative. Results indicate that participants were able to learn and use the temporal sequence to detect alarms. The large drop in performance when the temporal pattern changed provides support for our hypothesis that temporal sequence knowledge was used to monitor the gauges.

Since the performance results suggest that participants relied on temporal sequence knowledge for recognizing alarms, the next question is whether learning the temporal sequence was implicit or explicit. The questionnaire was designed to determine if participants could explicitly recognize the sequential patterns. Since participants did better than chance at correctly identifying which gauges were about to alarm, it would suggest that learning was explicit. However, performance on the task and performance on the posttest questionnaire were not correlated. The lack of correlation suggests that the explicit recall process required by the questionnaire was distinct from the presumably implicit temporal sequence learning process that facilitates alarm detection.

The current results are consistent with recent theories of reinforcement learning (Fu & Anderson, 2006; in press; Reichle & Laurent, 2006). In reinforcement learning, cue-action associations are strengthened by rewarding events (in this case, the sequence of alarms). In this task, repeated exposures to the same sequence of events reinforced the associations between alarms and implicitly guided the allocation of visual attention that maximized performance. As Reichle & Laurent show, the orchestration of successive eye movements often exhibit regularities that reflect the inherent patterns in the environment, and reinforcement learning seems to match well with the patterns of eye movements in reading tasks. Although we did not have moment-to-moment eye movement data to support this hypothesis directly in our task, the results suggest that participants were able to learn from trial-to-trial feedback based on the scores they were shown following each trial. Since performance improved across trials, it suggests that this post trial feedback provides enough support for their improvement in visual attention allocation to allow them to detect alarms more efficiently. However, further studies are required to better understand the role of reinforcement learning in visual sampling.

The current results have interesting implications to existing theories of skill acquisition. Many theories of skill acquisition propose that skill development starts with some form of general declarative representations of actions, and through experience, action selection and execution speeds up gradually (e.g., Anderson, 1982; Fitts & Posner, 1967). According to these theories, skill acquisition always starts with a slow, declarative stage,

where instructions are interpreted through verbal mediation, and working memory load is high because facts about the skill must be verbally rehearsed. Practice allows declarative knowledge to be slowly “compiled” into procedures that can be executed with little verbal mediation (e.g., Anderson, et al, 2004; Logan, 1988; Newell & Rosenbloom, 1981). Our results indicate that the utilization of external cues could be learned by the reinforcement learning mechanism that does not require the slow declarative stage that verbally rehearses the actions. In fact, we show that interactive skills involving actions that are cued by external information can be acquired directly and stored in forms that cannot be accessed verbally. Our results therefore suggest that existing models of skill acquisition may need to be extended to account for acquisition of interactive skill acquisition through some form of reinforcement learning, without necessarily going through the declarative-to-procedural process. As Ohlsson (1993) argued, it is likely that humans’ ability to learn preceded language, as it is too recent for special purpose brain mechanisms for verbal instructions to have evolved. Nevertheless, we believe that since learning to associate external cues to actions is an essential component of interactive skills, our results suggest that some form of reinforcement learning mechanisms should be included in future models of interactive skill acquisition.

These results have potential implications for improving the design of visual systems as well as training implications. Since this kind of temporal sequence learning is potentially implicit, the display could be re-designed to facilitate learning of temporal sequences by promoting the role of inherent patterns behind the events during training. Additionally, the implicit nature of learning suggests that it is difficult to train the observers to learn the sequences by verbal directions. Rather, training should be based on experiencing the sequence rather than verbal training.

Although scanning is adaptable to distal cues (e.g. patterns) when proximal cues (e.g. bandwidth) are not predictive, it is unclear whether observers will adapt to patterns when other cues are predictive. Future experiments will create situations where both the proximal cues are available and temporal patterns exist. We will attempt to determine if operators will adapt to more distal cues when more proximal, but less accurate cues are available.

ACKNOWLEDGEMENTS

The authors would like to thank Chris Bordeaux and Fernando Cordeiro for help with data collection.

REFERENCES

- Brunswik, E. (1952). *The Conceptual Framework of Psychology*. Chicago: University of Chicago Press.
- Cohen, A., Ivry, R., & Keele, S. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 17–30.
- Fu, W.-T., Anderson, J. R. (in press). Solving the Credit Assignment Problem: Explicit and Implicit Learning in a Probabilistic Sequential Choice Task. *Psychological Research*.
- Fu, W. & Anderson, J. (2006). From recurrent choice to skill learning: A model of reinforcement learning. *Journal of Experimental Psychology: General*, 135(2), 184-206.
- Holland, J. H. (1992). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Cambridge: MIT Press.
- Miller, S. M., Kirlik, A. & Kosorukoff, A. & M. D. Byrne (2004). *Ecological Validity as a Mediator of Visual Attention Allocation in Human-Machine Systems*. (AHFD-04-17/NASA-04-06). Savoy, IL: University of Illinois, Human Factors Division.
- Marcus, D., Karatekin, C., & Markiewicz, S. (2006). Oculomotor evidence of sequence learning on the serial reaction time task. *Memory and Cognition*, 34, 420-432.
- Moray, N. (1986). Monitoring Behavior and Supervisory Control. In K. Boff, I. Kaufmann and J. Beatty (Eds.), *Handbook of Perception and Human Performance*. New York, Wiley.
- Nissen, M., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19, 1–32.
- Reichle, E.D., & Laurent, P. A. (2006). Using Reinforcement Learning to Understand the Emergence of “Intelligent” Eye-Movement Behavior During Reading. *Psychological Review*, 113, 390-408.
- Senders, J. W. (1964). The human operator as a monitor and controller of multidegree of freedom systems." *IEEE Transactions on Human Factors in Electronics*.
- Senders, J. W. (1983). *Visual Scanning Processes*. Netherlands, University of Tilburg Press.
- Wickens, C. D., Goh, J., Helleberg, J, Horrey, W.J., and Talleur, D.A. (2003). Attentional models of multitask pilot performance using advanced display technology. *Human Factors* 45(3): 360-380.