

Visual Training Aids for Accelerating the Learning of Intuition

Aaron Novstrup, Sarah Lee, Terrance Goan

Stottler Henke Associates, Inc.

Seattle, WA

anovstrup@stottlerhenke.com, sarahlee@stottlerhenke.com, goan@stottlerhenke.com

ABSTRACT

Intuitive decision making (IDM) processes allow warfighters to synthesize great quantities of information at the speed required to act intelligently in highly complex, uncertain, and time-constrained tactical situations. A driver on combat patrol, for instance, does not have the time to consciously deliberate over the warning signs of an ambush, but must instead rely on the more immediate sense of intuition. Unfortunately, effective IDM only emerges through extensive experience, which for a warfighter operating in the real world can be both costly and life-threatening.

This paper reports on the effectiveness of perceptual augmentation for accelerating the acquisition of the situation recognition capabilities essential to intuition and for improving IDM. In two human studies (n=24 and n=60) involving a naturalistic, rapid situation categorization task (cf. Smith et al., 2017), we evaluated visual training aids designed to implicitly a) direct focus to the salient features of the situation, b) broaden attention to non-redundant perceptual inputs, and c) surface information from working memory that would otherwise be less accessible to IDM processes. Effects are measured by comparing subjects receiving visual assistance to those not.

Our findings suggest that it is possible to accelerate pattern learning by augmenting perception with stimuli that communicate information otherwise available only in working memory and by drawing attention to significant contextual cues. We also demonstrate that perceptual attention can be manipulated with subtle cues, apparently outside of conscious awareness—a capability that may prove essential in developing simulation-based implicit training methods that trigger the subconscious cognitive pathways hypothesized (Luu et al., 2010; Reber, 2013; Wan et al., 2011) to be associated with IDM. Finally, we find significant differences between intuitive and more deliberate subjects in how they allocate attentional resources (specifically eye gaze), raising the possibility that further study of the attentional behavior of rapid decision makers may lead to improved methods for cultivating effective IDM.

ABOUT THE AUTHORS

Aaron Novstrup is a software engineer and applied artificial intelligence researcher at Stottler Henke Associates, Inc., where he develops intelligent software systems with applications in information analysis, decision making, and training. Mr. Novstrup's current research focus is the application of machine learning and statistical modeling techniques to automated human state/performance assessment and monitoring. In ongoing work for DARPA and AFRL, he is developing a machine learning framework for real-time human state assessment based on input from unobtrusive sensors (e.g., eye gaze).

Sarah Lee is a software engineer at Stottler Henke Associates, Inc., with interests in user interface design and human-computer interaction (HCI).

Terrance Goan is a Senior Scientist at Stottler Henke Associates, Inc. At Stottler Henke Mr. Goan applies artificial intelligence and cognitive modeling to enhance human-computer interfaces for information analysis, decision making, and training applications.

Visual Training Aids for Accelerating the Learning of Intuition

Aaron Novstrup, Sarah Lee, Terrance Goan

Stottler Henke Associates, Inc.

Seattle, WA

anovstrup@stottlerhenke.com, sarahlee@stottlerhenke.com, goan@stottlerhenke.com

INTRODUCTION

Intuitive decision making (IDM) processes allow warfighters to synthesize great quantities of information at the speed required to act intelligently in highly complex, uncertain, and time-constrained tactical situations. A driver on combat patrol, for instance, does not have the time to consciously deliberate over the warning signs of an ambush, but must instead rely on the more immediate sense of intuition to recognize the danger and to act appropriately. Unfortunately, effective IDM only emerges through extensive experience, which for a warfighter operating in the real world can be both costly and life-threatening.

Research on the neurocognitive basis of intuition suggests that domain-specific experiences play a prominent role in its acquisition, implying that intuition is a product of associative learning that occurs outside of working memory and that permits information to be organized by preexisting knowledge on the time scale of milliseconds (Luu et al., 2010). A comparison of the brain activity of amateur and expert chess players presented with various chess positions and pressed to quickly come up with a move, for example, demonstrated that experts' brains had heightened levels of activity in areas of the brain responsible for visualization, episodic memory, and goal-directed behavior (Wan et al., 2011). *Implicit learning*, a form of learning that accounts for the dissociation between the ability to perform a task well and the ability to articulate (or make explicit) an understanding of the task, has been suggested as a possible mechanism for acquiring the knowledge structures involved in IDM (Reber, 2013). Situation recognition, the ability to recognize the similarity between a novel situation and one that has been encountered before, is also hypothesized to play a significant role in intuitive response (Klein, 1989).

Our work aims to demonstrate the potential for *augmented reality* and *immersive training* technologies to accelerate the acquisition of IDM skills and to improve the quality of rapid decisions. We focus here specifically on the potential of perceptual augmentation to accelerate the acquisition of situation recognition capabilities that are essential to intuition. In two human studies (n=24 and n=60) involving a naturalistic, rapid situation categorization task (cf. Smith et al., 2017), we evaluated two specific visual training aids. The first was designed to implicitly draw the attention of trainees to the cues essential to understanding situations encountered in immersive simulation environments. The second was designed to increase the perceptual accessibility of relevant cues (whether in training or in the field) by surfacing information that would otherwise be available only in working memory, where it is hypothesized to be less accessible to IDM processes.

EXPERIMENTAL TESTBED

Our experiments employed a naturalistic categorization task originally developed by Smith et al. (2017) to investigate the relationship between intuitive decision making and implicit learning in an operationally relevant context. The task requires participants to make a simple operational decision based on the features of a naturalistic perceptual stimulus (an animated 3D scene displayed on a computer monitor), varying along four primary dimensions: vegetation density, terrain topography, time of day, and weather conditions. The decision process is designed to be analogous to that which a small unit leader might undertake to select the appropriate formation to adopt in a squad on patrol. The subject must rapidly extract and integrate environmental information across perceptual features and then decide among a limited set of alternatives. The laboratory analog abstracts away from the operational model task in that a) no agents (whether friendly or hostile) appear in the environment, b) correct responses are mathematically determined and bear no resemblance to operationally relevant rules, and c) participants are not provided with explicit training of the underlying rules and instead must learn through trial-and-error with feedback. These abstractions facilitate experimental control over task difficulty, avoid performance variability arising from differences in prior knowledge, and permit investigation of implicit knowledge acquisition (Smith et al., 2017).

In this experimental framework, participants are shown a large sequence of unique scenes and given a brief period to choose one of three categories (e.g., Alpha, Bravo, or Charlie) for each scene. After a response is selected or time runs out, the participant is provided feedback indicating whether their response is correct or incorrect. Participants are provided with no description of the three categories and learn to recognize them only from the feedback provided.

The scenes, or “exemplars”, for each category are automatically generated based on a set of four-dimensional parameter vectors, with Unity-based scene generation software developed by Charles River Analytics. The four parameters correspond to the four primary factors of variation within the simulated environment (vegetation density, terrain topography, time of day, and weather conditions), and parameter values ranging continuously from 0.0 to 1.0 control scene construction by specifying a point along the corresponding environmental factor’s range of variation. The vegetation density parameter, for instance, determines the density of evergreen trees and shrubs in the generated scene, with values near 0.0 specifying sparse vegetation and values near 1.0 specifying dense vegetation. The four-dimensional parameter vectors are themselves generated by selecting points in a hyper-shell around each of three category “prototype” vectors.

Experimental design parameters control within- and between-category variation by determining the distances among the prototype vectors and the size of the four-dimensional shells from which the exemplar vectors are selected. The mathematical construction of the category prototypes and exemplars is designed to ensure that the categories can be distinguished only by integrating information across multiple dimensions. That is, the within- and between-category variation are carefully controlled so that the category decision boundaries appear indistinct in one- or two-dimensional projections.

EXPERIMENT 1: ATTENTION-GUIDING AUGMENTATION


Participants and Data Collection

In Experiment 1, 24 college students (13 females and 11 males, mean age = 24, SD = 5.3) were shown a sequence of 960 automatically generated stimuli over 5 blocks of testing, with 192 unique scenes per block. Participants were given up to 3 seconds to categorize each stimulus. The objective of this experiment was to validate our hypotheses that 1) a visual overlay could guide the visual attention of participants toward relevant stimuli in a pattern learning/recognition task 2) without being consciously recognized, 3) resulting in quantifiable performance gains 4) that persist even after the overlay is no longer available. Participants were randomized into two equally sized groups (test and control); participants in the test group were (sometimes) provided with an attention-guiding visual overlay, designed to subtly draw their attention to specific environmental cues that were strongly associated with each of the three categories. All participants reviewed the same 960 scenes, with the only difference being the visual overlay, visible only to test group subjects in a fraction of the scenes. The overlay was present in a decreasing fraction of the trials, eventually reaching zero for the entirety of the final block.

The environmental cues associated with each category in the original task described in (Smith et al., 2017) were distributed throughout each scene (e.g., time of day was indicated by the color of sunlight scattered by the sky and reflected by the earth and vegetation and by the length of any shadows cast by the sun) and therefore participants had no need to focus their attention on particular, *localized* cues in order to discriminate the categories. Operational tasks typically involve a mixture of diffuse cues, like those in Smith et al.’s original task, and more localized cues that intuitive decision makers must distinguish from many irrelevant stimuli. Indeed, a key aspect of intuition is the ability to allocate attention to the most important features of a situation and to ignore irrelevant ones.

Given our research focus on this attentional aspect of intuition, we added a localized environmental cue—specifically, one of three distinctive rock formations—to half of the generated scenes. The mathematical structure of the task was modified so that each of these rock formations was strongly associated with one of the three categories, identifying its associated category in 80% of the scenes in which it appeared. This strong, yet non-deterministic, association was intended to increase the discriminability of the categories in scenes in which such cues were present while avoiding conscious recognition of the cues. Thus, knowledge (whether implicit or explicit) of these hidden cues could theoretically boost a participant’s category recognition performance on trials in which they were present, especially for scenes near the category boundaries where the other environmental cues might provide insufficient discriminating power. Each of the rock formations had a recognizably distinct shape from the others but was otherwise unremarkable and appeared to be a natural part of the simulated landscape (see Figure 1). The formations were placed randomly near the center of the screen, with constraints on their apparent distance from the camera, and ray-casts were used to

ensure that they were not obscured by other objects (i.e., trees or hills). The rock formations were always oriented in the same way with respect to the camera to maximize the opportunity for participants to recognize their unique shapes.



**Figure 1. Top: Localized Cues associated with Alpha (left), Bravo (middle), and Charlie (right)
Bottom: Stimulus containing a Localized Cue**

The two groups of participants were treated identically, with the exception that participants in the test group were sometimes provided with an attention-guiding visual overlay. The overlay created a subtle fog effect designed to draw test group subjects' attention toward the localized environmental cues (i.e., the rock formations mentioned above). The overlay was virtually imperceptible in any static screenshot and was barely discernable even in the dynamic context of the experimental testbed. The number of trials in which the overlay appeared decreased in each testing block, until it was no longer present in the final block.

We gathered three types of data during the experiment: eye gaze data, capturing visual attention as participants completed the category learning task; task performance measures, including correctness of individual responses, category recognition accuracy over various temporal intervals, reaction time, response rate, and learning rate; and responses to a questionnaire administered to each subject after the final block of category learning, capturing any explicit knowledge the subject had formulated.

We captured participants' eye gaze data using a Tobii EyeX tracker at its sampling rate of over 80 Hz. In addition to logging the screen coordinates of each gaze point, we used 3D ray-casting to determine what kind of object in the scene (i.e., the sky, the ground, a tree, a rock) the subject was looking at for each sampled gaze point and logged this

information. We also logged an EyeX data stream identifying “fixations”—periods in which the eye remains relatively stationary.

From the logged gaze data, we derived a variety of summary statistics capturing each subject’s overall gaze behavior in each trial of the experiment. These statistics reflected the distribution of visual attention over different types of objects in the scene (e.g., fraction of attention devoted to the trees), over sub-regions of the screen (e.g., fraction of gaze in the top-left ninth of the screen), and over the scene as a whole (e.g., screen coordinate minima, maxima, variances), the overall rate of eye movement (i.e., the average distance between sampled gaze points per unit time), the distribution of fixated attention, and measures of gaze dynamics (e.g., the order in which different types of objects were fixated on).

We administered a debriefing questionnaire to all participants after their final block of testing. In addition to demographic questions on age, sex, level of education, and video game experience, the survey asked participants to describe “any patterns you noticed while playing the game, if any.” All participants provided detailed responses to this open-ended question.

Statistical Analysis

The eye gaze statistics were analyzed to understand differences in attentional behavior between various groups of participants (e.g., control vs. test), within individual participants at different points in time (e.g., before and after reaching above-chance performance), or across varying conditions (e.g., scenes for each category). We used two-way ANOVAs, controlling for learning effects across the five 192-trial blocks, to test for an effect of the intervention (the visual overlay) on the fraction of eye gaze allocated to localized cues (the special rock formations) in order to determine whether the intervention was effective at manipulating participants’ visual attention. More specifically, we tested for a between-subjects (control vs. test group) effect of the intervention, grouping by test block to control for learning, in a two-way ANOVA. We separately tested for a within-subject effect of the intervention (comparing test group trials with and without the attention-guiding overlay), also grouping by test block, in another two-way ANOVA. Finally, we analyzed the eye gaze behavior of the participants to determine whether “trained” participants exhibited distinctive attentional behaviors as compared to novices—behaviors that could potentially be exploited in attention-guiding visualizations in order to boost the learning of novices. To conduct these contrastive analyses, we employed two primary methods. First, we compared the distributions of eye gaze over the four main types of scenery (i.e., sky, ground/shrubs, trees, and rocks) using standard statistical tools such as boxplots, histograms, scatterplots, etc. Second, we attempted to train classifiers to distinguish the two groups of interest given only the features derived from their gaze behavior. The classifiers were trained on a fraction of the data and tested on the remaining data; classification accuracy on the test set serves as evidence of the distinctiveness of gaze behavior between two groups.

We tested for an effect of the intervention on subjects’ performance (in terms of accuracy within each block of testing) in a two-way ANOVA factored by experimental group and test block. We analyzed the association between attention to localized cues and task performance (in terms of categorization accuracy) by computing Pearson’s product-moment correlation between these variables. We also tested for a within-subjects effect of localized cues on performance in a two-factor ANOVA, grouping trials *with* versus *without* such cues and grouping by block.

Results

The intervention *was* effective at manipulating participants’ visual attention—the fraction of attention allocated to localized cues (i.e., special rock formations) exhibited both a statistically significant within-subjects difference depending on the presence of the attention-guiding overlay ($p < 0.0001$) and a statistically significant between-subjects difference across the test and control groups ($p < 0.00001$). These significant effects notwithstanding, *none of the participants mentioned the attention-guiding overlay* in their debriefing surveys, suggesting that it may have remained outside of their conscious awareness.

Somewhat surprisingly, however, despite both the effectiveness of the overlay at guiding attention and the strong association between each of the localized cues and the response categories, we found no significant effect of the intervention on the subjects’ performance. In fact, we observed a weak, but statistically significant, *negative* association between attention to localized cues and performance (Pearson’s $r = -0.325$; $p < 0.0003$). A within-subjects comparison of performance on trials with and without covert cues showed no significant effect, which suggests that the negative association may have been due to differences among subjects. For example, the negative association

might be explained by the fact that subjects with lower performance tended to focus near the center of the screen, where the covert cues were usually placed; thus, poor performers may have been more likely to look at the covert cues just by chance (combined with their tendency to take in less of the screen, noted below).

Taken together, these results suggest that the association between each of the rock formations and its correlated category may have been too subtle for participants to discern (implicitly or otherwise). That is, despite the strength of the association (with each rock formation identifying the correct category in 80% of the scenes in which it appeared), it may have been too difficult for subjects to recognize the significance of the *unique shapes* of the rocks. Certainly, none of the subjects mentioned the rocks in their verbal descriptions of the categories. In future work, it may be fruitful to repeat this experiment with localized environmental cues that vary on some other dimension (e.g., associating a different *type* of object with each of the three categories). Participants may also have been more likely to discover the associations if the localized cues had been non-natural objects (e.g., buildings, vehicles, people).

The participants in Experiment 1 generally formulated explicit decision rules, with the concreteness and correctness of these rules corresponding closely to their final performance in the task. (This result contrasts with Experiment 2, in which two distinct groups of top-performing subjects emerged exhibiting the standard characteristics of explicit and implicit learners.) Since our research goal was to shed light on *intuitive/implicit* rather than *deliberative/explicit* decision-making, these findings prompted several changes to the experimental design incorporated in the protocol for Experiment 2. In particular, we sought to encourage more *implicit* learning by: 1) reducing the response time limits to 2 seconds to allow less opportunity for conscious deliberation, 2) sampling exemplar parameter vectors from regions near the decision boundaries in 4D parameter space (which has the effect of increasing the amount of multi-dimensional evidence integration necessary for participants to achieve reliable category discrimination), and 3) developing tools/techniques for assessing sets of category exemplars *prior to* pilot testing by using their distribution in parameter space to predict the difficulty of the induced learning task and the degree of evidence integration required, facilitating optimization of the choice of category exemplars so as to balance the competing concerns of task difficulty and amenability to implicit learning.

By analyzing the allocations of visual attention over the three main types of objects in the scenes (i.e., sky, ground, and rocks), we found that top performers did indeed exhibit distinctive patterns of behavior as compared to novices. The two groups were easily distinguished by the fact that novices exhibited undifferentiated patterns of gaze behavior as compared to one another, while top performers each developed their own strategies for surveying the scenes. We hypothesized that with a sufficiently large sample, a handful of distinct groups of skilled performers would emerge, each with its own distinctive pattern of attentional behavior. This hypothesis is consistent with the multi-modal distributional patterns that we observed—the relatively small number of modes in the distributions of allocated visual attention indicate that participants adopted one of a handful of attentional allocation strategies.

EXPERIMENT 2: PERCEPTUAL SURFACING AUGMENTATION

Participants and Data Collection

In Experiment 2, 60 college students (31 females and 29 males, mean age = 24, SD = 4.4) were shown a sequence of 1,728 automatically generated stimuli over 9 blocks of testing, with 192 unique scenes per block. Participants were given up to 2 seconds to categorize each stimulus. The objective of this experiment was to validate the hypothesis that surfacing information from working memory could facilitate intuitive decision-making (IDM). Specifically, we hypothesized that since the cognitive mechanisms involved in IDM are believed to rely heavily on perceptual processing, augmenting perceptual inputs with information that would otherwise be available only in working memory would render that information more amenable to exploitation by these mechanisms and thereby accelerate the acquisition of intuitive decision-making (IDM) skills and/or improve rapid decision-making performance. We additionally designed the protocol to investigate the differences in eye gaze behavior between “intuitive” and “explicit” learners (as determined by the degree of evidence integration and explicit knowledge articulation exhibited by participants). Participants were again randomized into equally sized test and control groups. The groups were treated identically, except that the test group was provided with a visual overlay graphically depicting information that control group subjects needed to store in working memory.

The experimental task was again based on Smith et al.’s (2017) categorization task. In this version of the task, however, *the value of the weather parameter was not visually represented in the scene*. Instead, the weather parameter value was mapped to a temperature in degrees Fahrenheit ranging from 37°F to 98°F. It was held constant for sub-

sequences of up to 10 trials and was displayed to participants (in both groups) in textual form prior to each such subsequence. While control group subjects had to rely on working memory to store this information, test group subjects were additionally provided with a visual overlay that graphically depicted the “current” temperature throughout all trials (see Figure 2). Thus, only the test group was provided with a perceptual representation of the weather parameter value in close temporal proximity to response production.

We gathered the same data in Experiment 2 as in Experiment 1, except that the debriefing questionnaire was revised to allow for more detailed and structured responses. The revised questionnaire was similar to the one used in (Smith et al., 2017). It included open-ended questions intended to prompt participants to describe the response categories in their own terms as well as structured follow-up questions. The open-ended free response questions were: “Please describe your experience completing this task” and “Let’s say a good friend is coming in tomorrow to complete the same discrimination task you just completed and you want to give them a leg up. What would you tell them to ensure that they would be able to successfully discriminate between Alpha, Bravo, and Charlie?” Several follow-up questions asked participants to judge the importance of each of the environmental factors to the overall task and to qualify the relationship between each category and each environmental factor (e.g., “Could you tell me whether you thought Alpha had high (steeper), medium, or low (flatter) hills?”).

From the survey results, we computed an *integration score*, quantifying the degree to which subjects associated each category with multiple environmental factors (e.g., “Alpha was usually very hilly or densely forested”), and an *articulation score*, quantifying the degree to which subjects clearly and confidently expressed knowledge about the category structure. Vague descriptions or equivocations harmed the articulation score, while concrete descriptions increased it.




Figure 2. Stimulus with a Graphical “Temperature” Indicator (lower left)

Statistical Analysis

Participants were divided into groups based on the type and quality of their learning and these groups were compared with respect to their attentional behavior during the task. Consistent with the findings of Smith et al. (2017) in their similar experimental protocol, we found that participants fell into three general categories: “low performers”, whose performance plateaued without reaching a high (above-median) level of proficiency, “explicit decision-makers”, who explicitly formulated complex, deliberative rules to distinguish the categories, and “intuitive decision-makers”, who achieved high performance but were nonetheless unable to articulate sophisticated knowledge of the category

structure. Explicit decision-makers were formally distinguished based on their responses on our written questionnaire, using the integration and articulation scores described above—they articulated decision rules that required the integration of evidence across environmental dimensions. Although such evidence integration is a requisite to high performance on the task, intuitive decision makers achieved a high level of performance without articulating any conscious awareness of integrating evidence across dimensions.

Gaze statistics of the “explicit” and “intuitive” decision-makers in later blocks were averaged over 50-trial windows and the means of the two groups were compared with a two-sided t-test. To assess the discriminability of the two groups’ gaze behavior, we also trained a random forests classifier on 20% of the individual participants’ windowed (50-trial) gaze statistics for later blocks to classify each window according to whether it was generated by an explicit decision-maker or an intuitive decision-maker. The classifier was then tested on the held-out 80% of the 50-trial windows.

We quantified the learning rate and performance of each participant by fitting a linear mixed-effects model to participants’ rolling 20-trial accuracy. Time (in terms of the ordinal of each 20-trial window) was modeled as a fixed effect, and we included co-varying, normally distributed random effects of both subject and the subject-time interaction. Each participant’s learning rate was quantified by the slope of the linear model, estimating average increase in accuracy per unit time. Peak performance was quantified by the maximum accuracy estimated by the linear model over all of each participant’s 20-trial windows. Subjects who demonstrated no learning—those with an intercept under the linear model of less than 0.4 and a learning rate below 0.00004—were excluded from this analysis. We tested for an effect of the intervention (provision of the graphical temperature indicator to the test group) on learning rate with a two-sided t-test. A two-sided t-test was also used to compare the means of the test and control group subjects’ peak performance.

Results

The fitted linear learning model is shown in Figure 3. Under this model, we found that, after excluding four subjects who demonstrated no learning, the test group’s mean learning rate was 30% higher than that of the control group ($p < 0.034$, in a two-sided t-test). Subjects in the test group also attained higher peak performance on average, although this effect was not statistically significant (possibly due to a “ceiling” effect in this task). These findings lend credence to the hypothesis that intuitive decision making and the acquisition of intuitive decision-making skills might be improved by making relevant information available perceptually in cases in which that information would otherwise need to be stored in working memory. The results accord with the understanding of the intuitive system’s tendency to focus on information that is immediately available and of the strong connections among intuitive decision making, implicit learning/memory, and perception.




Figure 3. Learning Curves and Fitted Linear Model (by Subject)

29 participants achieved a block accuracy of at least 60% on their best block. Of these above-median performers, 16 were classified as explicit decision makers and 13 as intuitive decision makers. Analyzing their gaze statistics in later blocks over 50-trial windows, we found the group differences shown in Table 1. Qualitatively, the intuitive performers took in more of each scene (StdDevX/StdDevY) even though they did not move their eyes any faster (Eye Speed). They spent relatively more of their time than explicit performers looking at the sky and less looking at the trees and ground, although these differences may reflect the fact that covering more of each scene would result in a gaze distribution more closely matched to the distribution of the scene elements over the screen area. A random forests classifier trained on 20% of individuals’ windowed (50-trial) gaze statistics for later blocks was able to discriminate explicit and intuitive gaze behavior with over 78% accuracy (Cohen’s $\kappa = 0.53$; $p_{\text{binom}} = 0.0062$) on the held-out 80% of data. More work is needed to analyze the two groups on a more individualized basis and to determine whether subgroups emerge with distinctive patterns of attentional behavior.

Table 1. Differences in Mean Gaze Behavior between “Explicit” and “Intuitive” Decision-makers.

Gaze Statistic	Mean (explicit performers)	Mean (intuitive performers)	p-value
Sky¹	35.8	48.5	< 0.001
Trees¹	20.2	17.0	< 0.001
Ground¹	42.8	34.0	< 0.001
StdDevX²	3.64	5.09	< 0.001
StdDevY³	4.47	5.23	< 0.001
Eye Speed	0.396	0.394	0.81

DISCUSSION

The overarching goal of this effort has been to demonstrate the potential of augmented reality and immersive training technologies to accelerate the acquisition of intuitive decision-making skills and to improve the quality of rapid decisions. To this end, we investigated two specific forms of perceptual augmentation designed to aid in training/enhancing situation recognition skills. Experiment 1 tested an attention-guiding visual overlay designed to draw the attention of trainees to the critical features of an immersive simulation environment—the cues most essential to understanding situations encountered. Experiment 2 explored the effect of increasing the perceptual accessibility of relevant cues on situation recognition, using a visual overlay to surface information that would otherwise need to be stored in working memory.

Experiment 1 demonstrated the ability to manipulate visual attention with subtle perceptual cues, apparently outside of participants’ conscious awareness. The overlay did not appear to have any effect on category learning, however, perhaps due to issues in the experimental construction that prevented participants from making effective use of the localized cues to which their attention was drawn by the overlay. We suspect, in particular, that the link between these cues and the response categories was too subtle and that the localized cues offered little, if any, advantage over the redundant cues diffused throughout the environment. Increasing the “fuzziness” of the category discrimination boundaries in terms of the diffuse environmental cues, as we did in Experiment 2, might be expected to increase the relative importance of the localized cues. Replicating Experiment 1, with the experimental design improvements incorporated in Experiment 2, would thus provide a better measure of the effectiveness of our attention-guiding overlay.

The results of Experiment 2 indicate that it is possible to accelerate pattern learning and improve pattern recognition performance by providing perceptual stimuli that communicate information otherwise available only in working memory. This finding accords with theoretical models of intuitive decision-making that feature a prominent role for perception in the decision-making loop, and it suggests that augmentation technologies might enable warfighters to respond more intuitively in the field by making important contextual information more readily available in the form of perceptual inputs. Experiment 2 also demonstrated that explicit and intuitive decision-makers attend to perceptual cues in different ways and that skilled rapid decision-makers develop a handful of distinctive strategies for allocating attentional resources and quickly surveying the available evidence. Careful analysis of these differences and patterns may facilitate the development of training techniques or augmentation technologies that improve intuitive/rapid decision-making skills.

Future research is necessary to further clarify the mechanisms underlying intuitive decision-making and to provide new tools for accelerating the acquisition of IDM skills. Specific questions that warrant further investigation include:

¹ Percentage of time spent looking at the sky (or ground or trees).

² Standard deviation of the horizontal gaze coordinate, as a percentage of screen width.

³ Standard deviation of the vertical gaze coordinate, as a percentage of screen height.

1. **Can attentional behavior in operationally realistic scenarios be captured and analyzed to provide greater insight into the cues that skilled warfighters utilize to make rapid decisions?** Information regarding visual attention could feasibly be captured in live or simulated immersive training scenarios, using head-mounted gaze trackers / body-cams, virtual reality headsets, or screen-mounted gaze trackers like the Tobii EyeX, depending on the training environment. Recent advances in the field of computer vision, stemming from the development of highly capable “Deep Learning” neural networks, could facilitate the construction of neural networks for predicting skilled warfighters’ attentional behavior (i.e., networks which, given a series of images, predict which part of each image a warfighter will focus on). Furthermore, tools that have been developed to visualize the “reasoning” performed by such a neural network could yield generalizable insights into the warfighters’ behavior in rapid decision-making contexts.
2. **Can gains from attention-guiding graphical layers be decoupled from explicit knowledge acquisition?** Unforeseen issues with Experiment 1 prevented us from answering this question conclusively. First, all participants exhibited a high degree of explicit learning—a problem that we later remedied with changes to the construction of category exemplars. Second, despite the intervention working as intended to subconsciously draw participants’ attention to relevant cues, test group subjects were not able to derive any benefit (probably because it was too difficult to discern the association between those cues and the response categories). The success of Experiment 2 in bringing out a group of participants who exhibited the kind of intuitive decision-making often associated with implicit learning suggests that it may be beneficial to conduct a modified version of Experiment 1, incorporating an improved set of category exemplars (to increase the degree of implicit learning) and using environmental cues that participants more readily (but still implicitly) associate with the response categories. We hypothesize, for instance, that participants might more readily identify associations if the localized cues are encoded in the *types* of *non-natural* objects rather than in the *shapes* of *natural* objects.
3. **Can augmented reality be used to improve rapid decisions in real combat situations, by increasing the accessibility of key information to the warfighter’s “intuitive system”?** Experiment 2 demonstrated the benefits of surfacing critical information as perceptual input. Further research should be conducted to extend this result to more operationally relevant scenarios.

ACKNOWLEDGEMENTS

This material is based upon work supported by the U.S. Navy, Office of Naval Research, under Contract No. N00014-14-C-0025.

REFERENCES

- Klein, G. A. (1989). Recognition-primed decisions. In *Advances in man-machine systems research. Volume 5* (pp. 47–92).
- Luu, P., Geyer, A., Fidopiastis, C., Campbell, G., Wheeler, T., Cohn, J., & Tucker, D. M. (2010). Reentrant processing in intuitive perception. *PLoS ONE*, 5(3).
- Reber, P. J. (2013). The neural basis of implicit learning and memory: A review of neuropsychological and neuroimaging research. *Neuropsychologia*, 51(10), 2026–2042.
- Smith, M. K., Reuveni, B., Cohen, M. S., Grabowecky, M., & Reber, P. J. (2017). Developing a Naturalistic Categorization Task for Testing Intuitive Decision Making. In *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) 2017*.
- Wan, X., Nakatani, H., Ueno, K., Asamizuya, T., Cheng, K., & Tanaka, K. (2011). The neural basis of intuitive best next-move generation in board game experts. *Science*, 331(6015), 341–346.