

## **Developing a Naturalistic Categorization Task for Testing Intuitive Decision Making**

**Max Kailler Smith, Ben Reuveni,  
Michael S. Cohen, Marcia Grabowecy & Paul J. Reber  
Northwestern University  
Evanston, Illinois**

**max.k.smith@northwestern.edu, BenReuveni2020@u.northwestern.edu, michael.cohen@northwestern.edu,  
grabowecy@northwestern.edu, preber@northwestern.edu**

### **ABSTRACT**

Intuitive Decision Making (IDM) depends on knowledge that cannot be easily articulated. It does not reflect explicitly learned rules and guidelines. Rather, it is hypothesized to rely on implicit learning (IL). Basic science research on the phenomenon of IL provides a theoretical framework for understanding the acquisition of knowledge outside of conscious awareness from practical experience. This framework has the potential to accelerate the development of IDM during training and speed the acquisition of expertise. Here we describe a procedure and present a program of research based on adapting a more operationally relevant task to controlled laboratory conditions to bridge basic science and enhanced simulation-based training. The underlying task is one in which a complex decision is made based on environmental terrain characteristics, such as the formation in which to proceed with a patrolling infantry squad. This decision process is analogous to laboratory tasks in which participants learn to discriminate among a set of visual categories, but requires a new kind of task in which the visual stimuli are constructed from complex terrain dimensions. We defined a stimulus space based on four environmental dimensions: vegetation density, topography (hilliness), time of day and weather conditions. An artificial category structure was then defined within this stimulus space around three hidden prototypes. Participants learned these categories through trial-and-error with feedback about their decisions. Across three experiments, participants exhibited learning, increasing their decision accuracy across a range of task parameters selected to promote reliance on IL and use IDM. The resulting protocol will serve as a testbed for quantification of IDM effects and allow future work to examine training and educational interventions aimed at improving effective use of IDM. In addition, the task development process can serve as a model for bridging basic science research and operationally relevant domains.

### **ABOUT THE AUTHORS**

**Max Kailler Smith** is a research assistant in the Brain Imaging and Memory Lab at Northwestern University. He has been accepted to the Brain, Behavior and Cognition PhD program at Northwestern University and will begin Fall 2017.

**Ben Reuveni** is a graduate student entering his fourth year in the Brain, Behavior and Cognition PhD program at Northwestern University. His work in the Brain Imaging and Memory Lab focusses on memory systems interactions and category learning.

**Dr. Michael S. Cohen** is a postdoctoral researcher at Northwestern University in the Brain Imaging and Memory Lab. He is interested in how the value of information changes the way people learn, and in how people judge confidence in their own knowledge.

**Dr. Marcia Grabowecy** is a Research Associate Professor of Psychology at Northwestern University in the Brain, Behavior and Cognition Program. Dr. Grabowecy's research program is focused on multisensory perception and attention as fundamental scientific questions and in how multisensory integration can facilitate learning and decision making.

**Dr. Paul J. Reber** is a Professor of Psychology at Northwestern University and Director of the Brain, Behavior and Cognition program within the department. Dr. Reber's research examines the cognitive neuroscience of learning and memory and how these can be applied to enhance training, education and decision-making processes.

## **Developing a Naturalistic Categorization Task for Testing Intuitive Decision Making**

**Max Kailler Smith, Ben Reuveni,  
Michael S. Cohen, Marcia Grabowecky & Paul J. Reber  
Northwestern University  
Evanston, Illinois**

**max.k.smith@northwestern.edu, BenReuveni2020@u.northwestern.edu, michael.cohen@northwestern.edu,  
grabowecky@northwestern.edu, preber@northwestern.edu**

### **INTRODUCTION**

A vast array of operational tasks require individuals to use intuition to make decisions and act effectively, including, for example, actions undertaken by small unit leaders and information analysts. Currently, intuitive decision making (IDM) develops only incidentally after years of training and on-the-job experience, costing time, money, and potentially lives, as trainees are initially unable to effectively utilize IDM. We have undertaken a program of research aimed at leveraging the phenomenon of implicit learning to enable non-experts to be more effective decision makers, with a goal of accelerating training and improving decision-making accuracy and efficiency (Squire et al., 2014). To achieve this objective, we aim first to establish necessary scientific and technical methods for quantifying and improving IDM. Building on our history of laboratory research on the neurocognitive basis of learning and memory (Reber, 2013), we have begun to develop protocols for controlled studies in a more operationally relevant context. Here we describe construction of a new paradigm to characterize decision making in simulation-based training environments featuring implicitly-learned, covert statistical signals embedded in the stimuli.

Laboratory studies of human learning and memory have traditionally embedded a decision-making element via response required of experiment participants. As participants learn arbitrary visual categories, their knowledge is assessed by asking them to make decisions judging category membership for the experimental stimuli (Ashby & Maddox, 2005). This approach is used here to create an operationally relevant decision-making paradigm in which the role of intuitive decision making can be assessed and quantified. Whereas basic research laboratory studies have focused on simple, easily constructed, artificial stimuli, our new task incorporates a greater level of stimulus complexity in a mathematically defined space of stimulus dimensions that are rendered within a simulation-based environment.

In our novel protocol, participants make an operational decision based on environmental features of the visible terrain, defined in our simulation by varying vegetation density, terrain topology, time of day, and weather conditions. This process is analogous to a routine decision process that a small unit leader might have to make to select the appropriate formation to adopt for a patrolling squad. Environmental information would be extracted across perceptual features and then a decision would be made in order to direct further action (the formation to use). In the construction of a laboratory analog, our approach varies from the model task in several respects. First, to simplify the learning and decision process, no actors are included in the environment (neither friendly nor opposing) and the decision is entirely dependent on environmental characteristics. This aspect keeps the task simple enough to be learned within a typical laboratory session (1-2 hours) but allows for enough complexity to be challenging. Second, the correct responses in the laboratory task are entirely mathematically determined and bear no resemblance to rules that would be relevant in a simulated training environment. This provides maximal experimental control over task difficulty and avoids challenges in dealing with performance variability due to differences in prior knowledge. Third, no explicit training of the underlying rules is provided to participants, meaning that the correct decisions must be learned by experience. Our requirement to learn by experience is motivated by our research goal of examining the development and use of intuition in complex decision making (Nomura & Reber, 2012).

Across three experiments using our research protocol we show that participants can learn a novel categorization task based on exposure to simulated movement through simulated terrains. The task stimuli are 4 second movie clips showing artificial terrains generated with specific parameter values controlling the topography (hilliness), vegetation (tree/bush density), time of day (sunrise to sunset), and weather conditions (clear through mist/light rain). Within

roughly an hour of experience using trial-and-error with feedback, participants show significant improvement in their ability to categorize stimuli. Having established this testbed, our ongoing research will further develop this task as the basis for identifying measures of intuitive decision making via implicit learning.

At its core, this task is one that would be called a “prototype distortion” task in the cognitive psychology literature on category learning (Ashby & Maddox, 2005). Prototype distortion tasks use stimuli that are mathematically derived from a parameter-distortion algorithm applied to a set of underlying category prototypes (central tendencies) that are not shown to the participants. This approach is in contrast to rule-based category-learning tasks in which there is typically an explicit, discoverable rule that participants attempt to uncover during learning. Prototype tasks lend themselves well to implicit learning in which participants extract the statistically-determined similarity common across a set of novel stimuli and eventually come to be able to identify category membership without conscious retrieval of rules or prior exemplars. We defined a four-dimensional perceptual space that can be described by numeric values (between zero and one) on each of the four feature dimensions which were then rendered into a simulated terrain for participants to experience.

### Terrain Stimuli

Task stimuli were rendered by software designed and built by Charles River Analytics using the Unity platform. Input to the software was a configuration file that designated a ‘tuple’ of four numbers that determined the dimension values to be used in rendering a selected environment, together with a starting point and movement trajectory for the camera view. The four values controlled the “Hilliness”, “Vegetation”, “Time of Day”, and “Weather” during rendering. In particular, values close to 0 on the “Hilliness” dimension created uniform flat terrain maps, while values closer to 1.0 generated terrain maps containing cliffs and steep rocky hillsides. Values on the “Vegetation” dimension determined the overall density of evergreen trees and shrubs scattered throughout the terrain map. The “Time of Day” dimension ranged from noon to sunset. Lastly, low values on the “Weather” dimension produced calm environments, while decimal values closer to 1.0 generated cloudy, foggy, and windy conditions. The specified terrain was then rendered as a 4-second long video clip showing first-person perspective movement through the environment. Numerous stimuli were created per category to allow for practice without participants ever seeing the same video more than once. Videos were examined to exclude errors such as a physically impossible movement, for example, the camera moving accidentally and impossibly through a tree.



**Figure 1. Illustrations of 2 images of sample terrain stimuli for the category learning protocol that vary in vegetation density, hilliness, weather and time of day. Left image was rendered from tuple (0.1, 0.5, 0.2, 0.3) and right image was rendered from tuple (0.4, 0.4, 0.3, 0.5).**

### Category Prototypes and Exemplars

The embedded category structure was based on three prototypes (A, B, C) that were represented as equidistant points within the abstract (4-dimensional) feature space. In designing the categorization task, two general parameters allow for control of overall difficulty and can be expressed as within-category distance (how variable are the A’s from each other) and between-category distance (how different are the A’s from the B’s and C’s). For the experiments reported here, the between-category distance was set to 0.5 in the stimulus space (specific values are included in the details of each of the three experiments) that produced an appropriate level of difficulty for learning over a session of one to two hours. The within-category distance (0.140-0.141) was implemented by distorting specific stimulus examples away from the central prototype to create a sphere (or hypershell) of terrain stimulus

examples. This approach means that participants never saw the underlying prototypes during the task, which requires learning the categories from experience with feedback over trials.

We anticipated that the flexibility in defining the prototypes and stimulus exemplars would permit precise tuning of task difficulty. To study intuitive decision making, the task must be learnable, but also difficult enough so that participants do not immediately identify an explicit rule for category identification.

### Cover Story

To motivate learning, we provided participants with a cover story to engage them with the challenge of choosing actions that were appropriate for the three underlying categories. Participants were instructed that they were directing the exploration of an alien planet. On this planet, Alpha, Bravo, and Charlie travel methods for a supply operation are available. In response to a short video clip, participants were required to select the correct travel method to “arrive on time.” Thus, the task embedded a core decision process analogous to an operationally relevant process. However, there was no way in advance of experience with the stimuli to determine the correct answer, necessitating a trial-and-error approach to learn the correct travel method for the underlying terrain categories. After each response, feedback about the correct answer was provided. Decision-making accuracy was measured as percent correct over blocks of trials with the expectation that accuracy would start near chance (33%) and increase gradually as participants learned the underlying category structure.

### Participants

A total of fifty-four fluent English speakers between 18-35 years old were recruited from Northwestern University and the surrounding community. Participants were compensated \$15 per hour for their participation. Informed consent was provided by each participant prior to enrolling in the study. All procedures and recruitment materials were reviewed, approved and monitored by the Northwestern University Institutional Review Board. All recruited participants finished the experimental protocol. No participants were excluded from the experiments.

## EXPERIMENT 1: PRELIMINARY ASSESSMENT OF CATEGORY LEARNING

The participants’ learning of the three categories over a short protocol was evaluated to characterize their approach to the task and its difficulty using an initial set of values for between and within category distance.

### Procedure

Over 150 trials, participants (n=11) attempted to learn the correct travel method (Alpha, Bravo, Charlie) for generated terrain environments. With the first version of rendering software, terrain stimulus parameters were constrained to 10 steps from 0 to 1 and first-person perspective movement was limited to a straight line forward. The prototype values used to generate stimulus exemplars are listed in Table 1.

**Table 1. Experiment 1 Category Prototype Values**

	Hilliness	Vegetation	Time of Day	Weather
Alpha	0.3	0.4	0.3	0.3
Bravo	0.4	0.1	0.7	0.4
Charlie	0.1	0.3	0.5	0.7

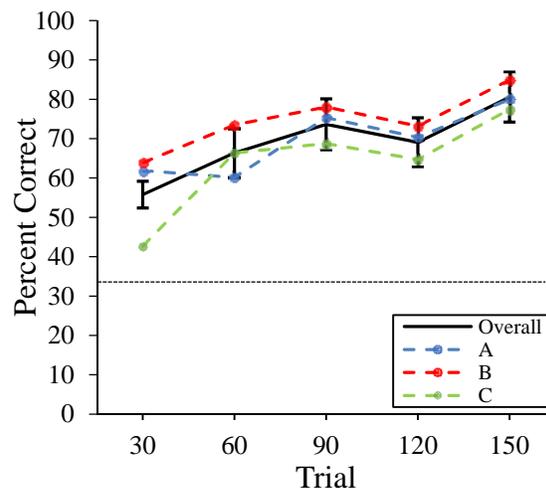
Participants read the task cover story and then attempted to discover the optimal travel method based on feedback provided after each trial. After making their category choice with a keypress that corresponded to one of the three methods of travel, a correct answer was followed by a green thumbs-up and a 500 Hz pure tone. After an incorrect response, participants saw a red thumbs-down, heard a 200 Hz tone and the saw the name of the correct navigation

method. Responses were limited to 2 seconds after the offset of the movie clip to promote rapid decision making. If no response was registered within the 2 seconds, a screen would appear requesting that the participant respond faster on the next trial. Trials were pseudo-randomized in 30 trial blocks (10 trials per category per block) so that trials belonging to a single category could not occur more than three times in a row. This type of pseudo-randomization by block ensures each category is represented equally in each block as well as mitigates the potential confound of superstitious responding to spurious patterns in the trial order. This superstitious responding behavior is akin to the “gambler’s fallacy” in that multiple occurrences of a single type of trial in a row can mistakenly lead participants to believe that type of trial is less likely to occur next, even though all trial types have equal probability of appearing next (Tversky & Kahneman, 1971). Participants were prompted to take a short break after every 50 trials to avoid fatigue and keep the entire protocol to approximately an hour.

After the category-learning task was complete, a short, open-ended post-session interview was administered in which participants were asked to describe their experience performing the task. Participant reports of strategies, identified patterns, and any other comments about the task were reviewed and summarized qualitatively.

## Results

Participants generally learned all categories at greater than chance level, achieving an average of 80.5% (SD 21.0%) accuracy during the final task block (Figure 2). Final block performance was significantly greater than chance (33%),  $t(10) = 7.44$ ,  $p < 0.001$ , indicating successful learning. In addition, a paired samples t-test between the final block and first block (55.8% correct, SD 11.2%) showed performance significantly improved with practice,  $t(10) = 3.66$ ,  $p < .05$ . However, accuracy on the first block of 30 trials was already substantially above chance performance,  $t(10) = 6.60$ ,  $p < 0.001$ . In addition, a two-way ANOVA was conducted to examine possible differences across categories and the effect of practice (first half of trials vs last half of trials). There was a significant category effect,  $F(2, 20) = 4.34$ ,  $p < 0.05$  indicating that not all categories were discriminated at the same rate. There was a marginal effect of practice across blocks,  $F(1, 10) = 4.50$ ,  $p = 0.059$ , and no interaction between category and learning,  $F(2, 20) = 0.28$ ,  $p = 0.75$ . A planned pairwise comparison using t-tests revealed that the category difference was due to lower performance on Charlie compared to Alpha,  $t(10) = -2.37$ ,  $p < 0.05$ , and Charlie compared to Bravo,  $t(10) = -2.49$ ,  $p < 0.05$ . Alpha and Bravo did not differ in performance,  $t(10) = 1.22$ ,  $p = 0.24$ .



**Figure 2. Overall and category accuracy in Experiment 1. Each point is the average accuracy for 30 trials. Error bars represent standard error of the mean.**

In post-session interviews, many participants reported that the weather dimension was particularly distinctive for the Charlie category stimuli in that it was the only one to include salient rain, which was only present in stimuli greater than or equal to 0.6 on the weather dimension. This explicit rule (rain = Charlie) was discovered very rapidly, leading to moderately high accuracy for stimuli from the Charlie category. Accuracy was not at ceiling because not all Charlie stimuli included rain. However, Experiment 1 verified that the experimental framework using an arbitrary ‘supply method’ categorical decision was feasible. To review, those decisions were based on mathematically structured characteristics of terrain environments that were rendered into a first-person perspective simulation. To reduce reliance on the salient weather dimension, the simulation rendering software was adjusted for subsequent experiments to eliminate the salient rain and make the weather dimension steps smaller. Anticipating that this change would make the task more difficult, and more appropriate for implicit learning, the number of trials was also increased for Experiment 2.

## EXPERIMENT 2: INCREASED NUMBER OF TRIALS AND FINER STIMULUS CONTROL

For Experiment 2, in addition to removing stimuli with rain, the software was upgraded to allow more fine-grained steps on each dimension and new features for control over the first-person camera trajectory. A new set of 300 stimuli for the learning protocol was re-rendered using these new features.

### Procedure

Nineteen new participants were recruited and consented to participate in Experiment 2. The category prototype stimulus values used to create the exemplars for this experiment are listed in Table 2. The path of movement through the simulated environment was changed from a straight line to a backwards “S” in order to provide more opportunities to take in features of the environment while participants observed the video clip.

**Table 2. Experiment 2 Category Prototype Values (rounded to thousandths decimal place)**

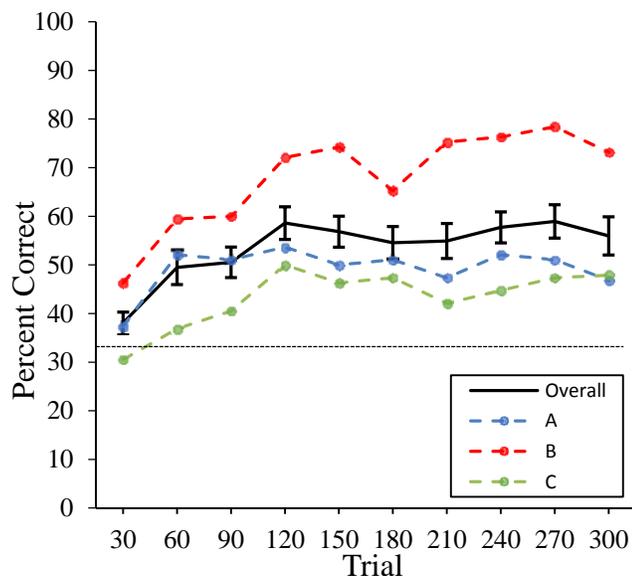
	Hilliness	Vegetation	Time of Day	Weather
Alpha	0.124	0.110	0.148	0.307
Bravo	0.377	0.471	0.079	0.531
Charlie	0.580	0.209	0.313	0.239

The instructions and trial structure was similar to that of Experiment 1 but included 300 trials with short breaks every 50 trials. After the category-learning task was complete, a post-session structured interview was administered to assess explicit knowledge of the three categories. The post-session interview included questions about the rules (heuristics) used by participants and posed a hypothetical situation in which participants were asked to coach another participant on how to do the task. The participant was then informed about the four underlying feature dimensions, “Hilliness”, “Vegetation”, “Time of Day”, and “Weather” and was asked to rate on a 1 to 10 scale the importance of each dimension and whether each of the categories was best characterized by low, medium or high values on each feature dimension.

As in Experiment 1, after the category-learning task was complete, a short, open-ended post-session interview was once again administered. Participant reports of strategies, etc., were reviewed and summarized qualitatively.

### Results

Over the 300 trials of learning, participants’ accuracy increased from 38.1% correct (SD 9.6%) on the first 30 trials to 55.9% correct (SD 17.1%) on the final 30 trials (see Figure 3). Performance on the final block was again better than chance (33%),  $t(18) = 5.76$ ,  $p < 0.001$ , reflecting successful learning. For this experiment, learning began at levels closer to chance performance and increased more gradually than it did in Experiment 1 with the salient weather rule no longer available to guide decision-making. A paired-samples  $t$ -test showed that performance in the final block was significantly better than the first block,  $t(18) = 4.68$ ,  $p < .001$ , indicating that practice enhanced



**Figure 3. Overall and category accuracy in Experiment 2. Each point is the average accuracy for 30 trials. Error bars represent standard error of the mean.**

performance. However, examination of the learning rate for the three categories revealed that there was a much larger increase in performance for the Bravo category than for Alpha or Charlie. A two-way ANOVA was conducted to examine the effects of practice (first half of trials vs. last half of trials) and category on performance. This revealed a significant practice effect,  $F(1, 18) = 10.22, p < 0.001$ , and category effects,  $F(2, 36) = 30.14, p < .001$ , as well as a reliable interaction between practice and category,  $F(2, 36) = 4.31, p < 0.01$ . A planned pairwise comparison using t-tests of category showed that this effect was due to significantly better learning for the Bravo category compared to Alpha,  $t(18) = 4.17, p < 0.001$ , and Charlie,  $t(18) = 5.47, p < 0.001$ . There was no significant difference in the effect between Alpha and Charlie,  $t(18) = 1.31, p = 0.19$ .

The significantly better performance on the Bravo category suggested that participants were only discriminating this category. Post-session interviews revealed that 11 out of 19 participants used the weather dimension (presence of fog) to easily distinguish Bravo from Alpha and Charlie. Of note, a single category strategy can produce apparently above-chance performance even on the unlearned categories by making those decisions (Alpha or Charlie, here) a 50/50 guess, which will appear to be better than the 3-alternative chance rate (33%). Thus, the set of prototypes used for Experiment 2 appeared to lead to modest learning overall through a focus on a single, discriminable category (Bravo).

### EXPERIMENT 3: BALANCED LEARNING ACROSS CATEGORIES

For Experiment 3, a new stimulus and prototype set was constructed that was intended to more evenly balance discriminability across the three categories. It was expected that this would reduce the overall difficulty of the task and elicit better learning within the one-hour experimental protocol. In addition, the rate at which participants were able to make their decisions indicated that we could further increase the number of trials in the experiment to examine a more extended learning curve.

#### Procedure

Twenty-four new participants were recruited and consented to participate in Experiment 3. Participants experienced 450 trials with breaks every 50 trials. Category prototypes used to generate exemplars are listed in Table 3. Trial order and experimental protocol were otherwise identical to Experiment 2.

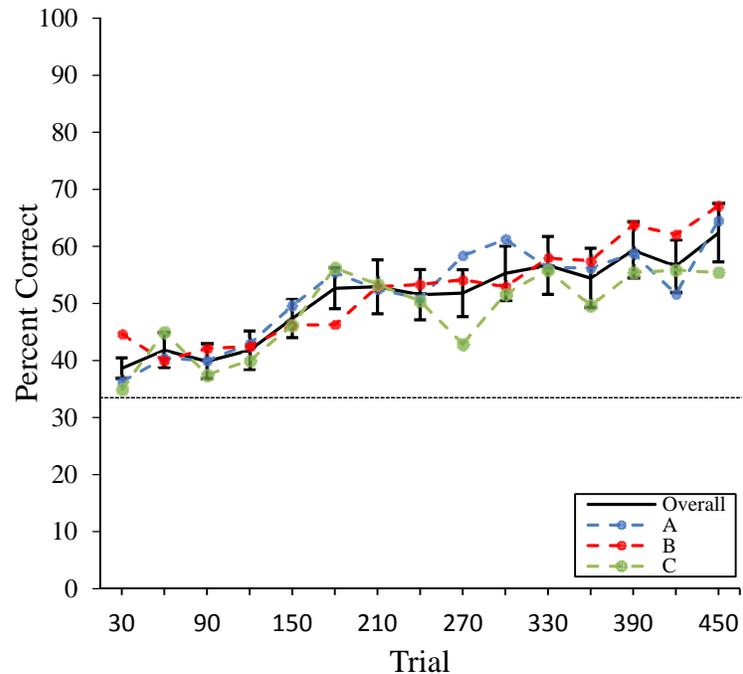
For this experiment a new, structured post-session interview was developed to allow for a more quantitative evaluation of participants' reliance on explicit, reportable categorization strategies. In addition to open-ended questions about general strategy use, participants were further encouraged to provide whatever guidance they could that might be helpful to another person in the experiment (i.e., what advice would they give a friend). Participants also rated the importance of each dimension for each category and rated their confidence in their own knowledge on 1-10 Likert scales. The responses of the participants on these questions was scored by four raters who were blind to the experimental hypothesis in order to avoid any bias in the evaluation of explicit, strategic knowledge reported.

**Table 3. Experiment 3 Category Prototype Values**

	Hilliness	Vegetation	Time of Day	Weather
Alpha	0.891	0.887	0.390	0.582
Bravo	0.737	0.912	0.701	0.224
Charlie	0.476	0.857	0.290	0.324

## Results

Participants demonstrated learning in the categorization task (Figure 4), exhibiting gradual increases in decision accuracy across the learning session and performing similarly for the three underlying categories. Performance on the first block appeared close to chance (38.6% correct,  $SD = 9.05\%$ ), however a two tailed t-test showed that participants were already learning the categories and even the first 30 trials were performed significantly better than chance,  $t(23) = 2.85$ ,  $p < 0.01$ . On the final block of the experiment, participants were accurate on 62.3% ( $SD 24.9\%$ ) of trials. A two-tailed t-test confirmed that final block performance was significantly above chance,  $t(23) = 5.71$ ,  $p < .001$ . In contrast to Experiment 2, a two-way ANOVA of practice (first half vs last half of total trials) and category did not show a main effect of category,  $F(2, 46) = 0.97$ ,  $p = 0.38$ , nor an interaction between category and practice,  $F(2, 46) = 1.11$ ,  $p = 0.33$ , signifying more even learning of the three categories. However, there was a significant practice effect, as expected,  $F(1, 23) = 14.18$ ,  $p < 0.01$  reflecting robust and more uniform learning of the three categories. Additionally, a paired samples t-test revealed a significant difference between performance in the first and last block of trials,  $t(23) = 5.17$ ,  $p < 0.001$ , indicating that performance improved reliably with practice.



**Figure 4. Overall and category accuracy in Experiment 3. Each point is the average accuracy for 30 trials. Error bars represent standard error of the mean.**

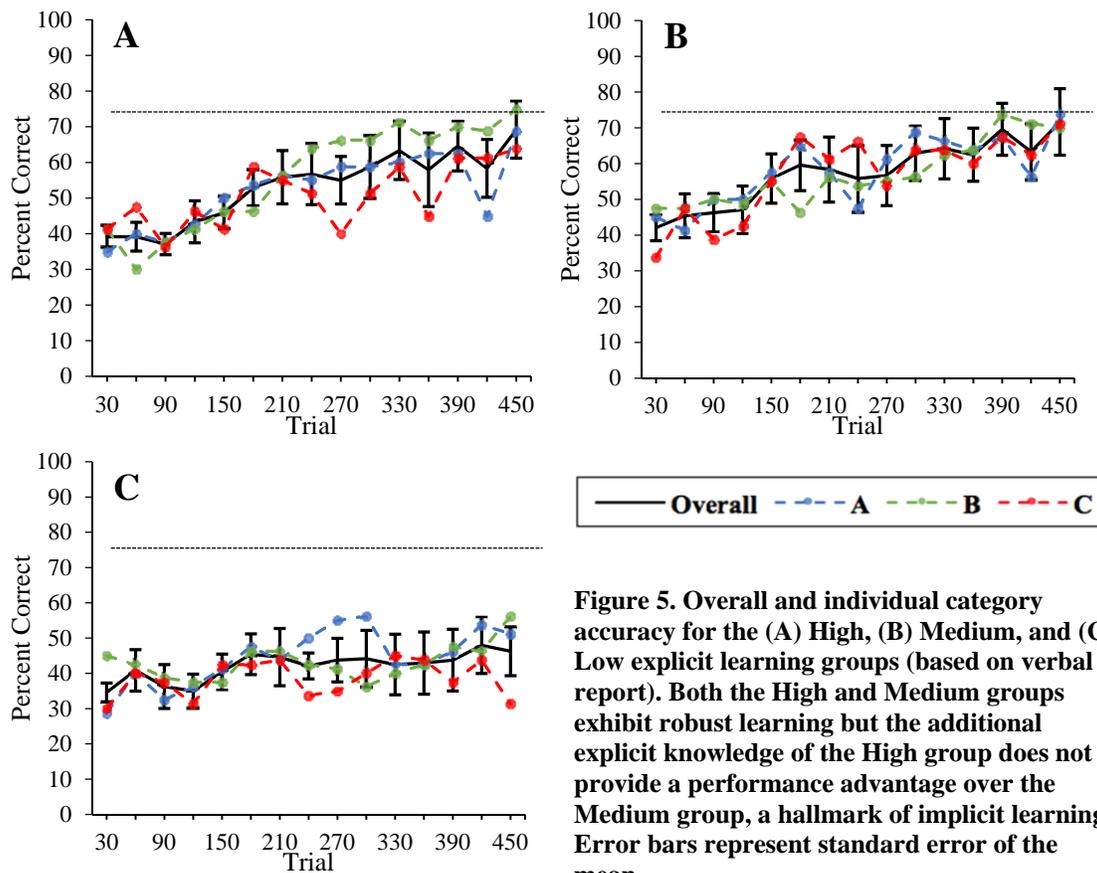
On post-session interviews participants reported a variety of potential rules that they described as guiding their decision-making process. While some participants reported general rules relating to the actual stimulus parameters, other participants constructed idiosyncratic and often inaccurate rules. For instance, always choosing Bravo if a rainbow was present, or distinguishing Alpha and Charlie using purple or white flowers (these features were not reliable predictors). To quantitatively evaluate the post-learning interviews, a scoring rubric was developed for the post-session interviews providing a numerical assessment of the importance of each dimension as verbally reported by each participant. Responses were evaluated on two scores. An 'articulation' score reflected a clear and accurate statement about how the dimensions related to the learned categories, scored on a 1-10 scale. In addition, a 'conjunctive' score counted accurate statements about the role of each dimension for each category with a maximum score of 12 (3 categories, 4 dimensions). The 'conjunctive' score was averaged across both the verbal report of the strategy used by the participant and a second interview question asking the participants to give advice to somebody else who was going to complete the experiment. Four naïve raters applied the scoring rubric to each participant's interview data and the raters' scoring was highly consistent (Intraclass Correlation Coefficient,  $ICC = .92$  for conjunctive scores;  $ICC = .74$  for articulation).

Based on the interview responses, participants were grouped into three levels based on the amount of explicit knowledge expressed verbally after completing the experiment. A 'Low' explicit knowledge group included participants who scored less than 5 on the articulation score or less than 3 on the conjunctive score. A 'Moderate' explicit knowledge group scored less than 6 on the articulation score and less than 5 on the conjunctive score. The 'High' explicit knowledge group scored greater than 6 on articulation score or greater than 5 on conjunction score. Average ratings and categorization accuracy of each group is shown in Table 4 (see also Figure 5). Participants with more explicit knowledge generally did better at categorization ( $r = 0.52$  for articulation,  $r = 0.40$  for conjunction) indicating that explicit knowledge contributed at least somewhat to performance. However, a number of participants performed well at the categorization task (e.g., within the Moderate group) without being able to articulate much

information about the category structure, indicating an important contribution of implicit learning.

**Table 4. Explicit knowledge reported in post-learning structured interview** (numbers in parentheses reflect SD of the mean)

Explicit Knowledge Groups	Avg. Articulation Score (0-10)	Avg Conjunction Score (0-12)	Average % Correct (Final 30 trials)
High (n = 8)	3.375 (1.10)	1.43 (1.61)	45.20 (22.23)
Medium (n = 8)	5.56 (0.25)	3.59 (0.53)	66.76 (21.87)
Low (n = 8)	6.40 (0.42)	5.65 (1.02)	62.49 (22.55)



**Figure 5. Overall and individual category accuracy for the (A) High, (B) Medium, and (C) Low explicit learning groups (based on verbal report). Both the High and Medium groups exhibit robust learning but the additional explicit knowledge of the High group does not provide a performance advantage over the Medium group, a hallmark of implicit learning. Error bars represent standard error of the mean.**

## **GENERAL DISCUSSION**

Using the literature on cognitive psychology of category learning as a reference, we constructed a novel task to study decision-making based on the characteristics of terrain in simulated environments. In a single one-hour session, community participants (undergraduate) learned to distinguish among three artificially defined, covert categories of terrain. Improved categorization performance was observed in all three Experiments, although results demonstrated the challenge of balancing between-category difficulty and dimension saliency across different task variants. Experiment 1 demonstrated that participants tend to focus on a particularly salient value within a single dimension if available. In Experiment 2, one of the embedded categories was more easily discovered, leading to a partial learning strategy. In Experiment 3, successful gradual learning of all three of the embedded categories was observed over 450 trials, thus providing a testbed in which decision-making accuracy could be assessed through the trial-and-error learning process.

For our purposes of examining intuitive decision making, the responses in the post-session structured interview indicating some participants were aware of the hidden category structure pose a challenge. Our goal is to define a decision-making context in which participants report using a hunch (intuition or instinct) and yet perform reliably better than chance, a phenomenon seen in many implicit learning tasks (Reber, 2013). Approximately a third of the participants in Experiment 3 exhibited a profile consistent with implicit learning (the Partial group) in which relatively high levels of decision accuracy occurred with low explicit knowledge. Future work will aim to improve the rate of intuitive decision making in this task through increasing the reliance on implicit learning. Specifically, covert auditory cues that are predictive of the correct decisions will be added probabilistically to some decision trials. We predict that these cues will influence decision making via intuition, evoking a hunch about the correct answer when the cues are available and allowing us to quantitatively measure intuitive decision making by comparing performance when these additional cues are or are not available.

The information learned by participants over the experimental session was mathematically constructed to be arbitrary and not connected to operationally relevant decisions about moving across an unknown environment. By simplifying the decision process, our results demonstrate how a simulation-based environment can be used to examine learning over a large number of trials in a short period of time (450 trials in one hour). With a set of just three options (A/B/C) made in rapid succession, it was possible to assess performance (as percent correct accuracy) and the evolution of learning. The approach described raises the possibility that this style of simulation-based training may be used to create more operationally relevant task stimuli that can still benefit from basic science in learning and decision-making. Controlled experiments can then provide research tools to quantify, and methods to optimize, training protocols.

The method described here for creating a new task protocol suitable for experimental manipulation could be applied and extended to a wide variety of operationally relevant training domains. The key features of the process are a parameterized characterization of the underlying space of training stimuli and tools for rendering the stimulus space into a visualization (here, the terrain environment). Techniques of experimental psychology can then be used to structure controlled conditions to isolate theoretically important features of the learning process and quantitatively assess the impact of controlled manipulations of the training context. This approach opens up a broad new set of possibilities for applied research into improving protocols for simulation-based training building directly on research from basic laboratory research.

## **ACKNOWLEDGEMENTS**

The reported results were supported by a grant from the Office of Naval Research (N00014-16-1-2251) to support research examining the role of implicit learning on the development of intuitive decision making.

## **REFERENCES**

- Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annu. Rev. Psychol.*, *56*, 149-178.
- Nomura, E. M., & Reber, P. J. (2012). Combining computational modeling and neuroimaging to examine multiple category learning systems in the brain. *Brain Sciences*, *2*(2), 176-202.
- Reber, P.J. (2013). The neural basis of implicit learning and memory: A review of neuropsychological and neuroimaging research. *Neuropsychologia*, *51*, 2026-2042.
- Squire, P., Cohn, J., Nicholson, D., Nolan M., Reber, P.J., Oudiette, D, Niehaus, J., Geyer, A., & O'Neill, L. (2014). Towards enhancing intuitive decision making through implicit learning. Interservice/Industry Training, Simulation, and Education Conference (IITSEC) 2014.
- Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological bulletin*, *76*(2), 105.