UNIVERSITY OF CALIFORNIA, SAN DIEGO

# Incidental Learning and Explicit Knowledge

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Psychology

by

Randy Tran

Committee in charge:

Professor Harold Pashler, Chair
Professor Andrea Chiba
Professor Gary Cottrell
Professor Timothy Rickard
Professor Edward Vul

2017

The Dissertation of Randy Tran is approved, and it is acceptable

in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2017

## DEDICATION

To my family and friends,

the love of my life, Isabel,

and Caspurr.

## TABLE OF CONTENTS

## LIST OF FIGURES

# LIST OF TABLES

## ACKNOWLEDGEMENTS

First, I would like to acknowledge Liz Schotter, Sarah D'Angelo, Bernhard Angele, the former members of the Rayner Lab, and Keith Rayner, for without their support and encouragement, I would not have pursued graduate school.

Noriko, Jonas, and other members of the Learning, Attention, and Perception Lab deserve acknowledgements for their assistance throughout my graduate career from comments on experimental design to running subjects. They certainly added color to my day-to-day graduate life.

My last and most appreciative acknowledgement and sincerest thanks goes to my advisor, Hal Pashler. His support, advice, and constructive criticism has taught me a great deal on how to become an independent researcher and how invaluable it is to have a variety of skills under my belt.

Chapter 1, in full, is a reprint of the material as it appears in the *Royal Society Open Science* 2017. Tran, Randy; Vul, Edward; & Pashler, Harold. The dissertation author was the principal researcher and author of this material.

Chapter 2, in full, is a reprint of the material as it appears in *PLOS ONE* 2017. Tran, Randy; Pashler, Harold. The dissertation author was the principal researcher and author of this material.

Chapter 3, in part, is currently being prepared for submission for publication of the material. Tran, Randy; Pashler, Harold. The dissertation author was the principal researcher and author of this material.

## VITA

| | |
|---|---|
| 2017 | Ph.D. in Psychology, University of California, San Diego |
| 2012-2017 | Graduate teaching assistant, University of California, San Diego |
| 2011-2012 | Undergraduate teaching assistant, University of California, San Diego |
| 2013 | M.A. in Psychology, University of California, San Diego |
| 2012 | B.S. in Human Biology & B.S. in Psychology, University of California, San Diego |

## PUBLICATIONS

Randy Tran, Edward Vul, and Harold E Pashler, "How effective is incidental learning of the shape of probability distributions?", *Royal Society Open Science*, 2017.

Cory A. Rieth, Ben D. Amsel, Randy Tran, and Maia B. Cook, "Issues and Advances in Anomaly Detection Evaluation for Joint Human-Automated Systems", *Advances in Human Factors in Robots and Unmanned Systems*, 2017.

Randy Tran and Harold E Pashler, "Learning to exploit a hidden predictor in skill acquisition. Tight linkage to conscious awareness", *PLOS ONE*, 2017.

R. Dirk Beer, Cory A. Rieth, Randy Tran, and Maia B. Cook, "Framework for Multi-Human Multi-Robot Interaction. Impact of Operational Context and Team Configuration on Interaction Task Demands", *AAAI Spring Symposium Series*, 2017.

Randy Tran, Doug Rohrer, and Harold E Pashler, "Retrieval practice. the lack of transfer to deductive inferences", *Psychonomic Bulletin & Review*, 2015.

Elizabeth R Schotter, Randy Tran, and Keith Rayner, "Don't believe what you read (only once) comprehension is supported by regressions during reading", *Psychological Science*, 2014.

Bernhard Angele, Randy Tran, and Keith Rayner, "Parafoveal–foveal overlap can facilitate ongoing word identification during reading. Evidence from eye movements. *Journal of Experimental Psychology. Human Perception and Performance*, 2013.

## ABSTRACT OF THE DISSERTATION

Incidental Learning and Explicit Knowledge

by

Randy Tran

Doctor of Philosophy in Psychology

University of California, San Diego, 2017

Professor Harold Pashler, Chair

In skill acquisition, use of predictive relationships are needed to perform optimally and are assumed to be acquired independently of awareness. In three chapters, I investigated how individuals exploit the predictive relationships to their advantage and address the following questions: 1) Does exploitation of predictive relationships occur when there is no explicit knowledge of the underlying structure? 2) Might some of the findings in implicit learning have very limited generalizability? 3) Lastly, does averaging across subject data mask what is learned by individuals? Altogether, my dissertation

revealed people 1) strongly favored simple and verbalizable relationships, 2) used explicit knowledge of the predictive task-relevant features for optimal performance, 3) used very different strategies that were only revealed in post-experiment questioning.

## Introduction

In acquiring a skill, it is often important for people to find ways to exploit predictive relationships that may exist in the task domain. These relationships may be simple but they can also be quite complex. My dissertation examines unresolved questions about how this exploitation process occurs.

Some researchers have suggested that people utilize an explicit hypothesis-testing process for predicting simple relationships when the predictive features are easy to verbalize (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998). On the other hand, they propose an implicit and typically nonconscious procedural learning process is used to detect and utilize complex relationships when the predictive features are difficult or impossible to verbalize (e.g., Ashby & Maddox, 2005; Reber, 1989; Sanchez, Gobel, & Reber, 2010). The weather prediction task (Knowlton, Squire, & Gluck, 1994) is a typical example used to argue for the dissociation between the two learning systems. In this task, amnesiacs and healthy controls had to guess a weather outcome (sunshine or rain) when presented with a combination of cue cards. The true weather outcomes were probabilistic in nature thus making the predictive features difficult to verbalize. Experimental analysis of the task has showcased implicit learning (i.e., amnesiacs' weather predictions were comparable to control subjects).

However, a resurgent interest in awareness during learning has spurred research that seems to indicate that even in tasks that require integration of probabilistic outcomes, conscious explicit hypothesis testing may take place (e.g., Haider & Frensch, 2005; Lagnado, Newell, Kahan, & Shanks, 2006; Price, 2009; see also, Knowlton, Squire, & Gluck, 1994).

The main questions posed in my dissertation research are: 1) Does exploitation of predictive relationships occur when there is no explicit knowledge of the underlying structure? 2) Might some of the findings in implicit learning have very limited generalizability? 3) Lastly, does averaging across subject data mask what is learned by individuals? In three chapters, a common theme will emerge. Explicit knowledge plays a bigger role in incidental learning than previously thought.

In Chapter 1, subjects were exposed to a unidimensional-bimodal distribution (i.e., the spatial location where dots appeared on a line). Incidental learning was then assessed by asking subjects to generate novel samples. The results suggest participants do not implicitly learn the underlying complex distribution (i.e., bimodality) and default to a simple distribution (e.g., uniform). Only when the complex distribution was discretized (i.e., bimodal distribution is made extremely apparent) did participants show incidental learning of the underlying structure.

Chapter 2 expands on these findings by examining whether subjects require conscious awareness to exploit predictive features. The two experiments employed relevant predictive features that were binary and deterministic that were among varying irrelevant features. Across both experiments, subjects that were unable to verbalize the rule did not show any implicit learning. In other words, explicit knowledge and ability to verbalize the rule were necessary to exploit the predictive features to subjects' advantage. However, an alternative explanation could be the usage of a single deterministic dimension begs for an all-or-none process (e.g., Sergent & Dehaene, 2004; Smith & Kounios, 1996).

Therefore, the set of experiments in Chapter 3 were designed to have subjects learn an orthogonal 2-feature category space (size and hue) constructed from two uniform distributions. This approach creates a complex decision boundary that requires a conjunction of two dimensions. The findings from Chapter 3 converge with results from Chapters 1 and 2 where numerous subjects default to simple structures (i.e., using only one dimension) and only subjects that were able to verbalize both size and hue were able to classify exemplars with the correct diagonal boundary.

The findings from the three chapters suggest people favor simple and verbalizable relationships and use explicit knowledge of predictive features for optimal performance. Implications for each chapter will be discussed in their respective manuscripts.

**References**


Ashby, F. G., Alfonso-Reese, L. A., Turken, U, & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*(3), 442–481. http.//doi.org/http.//dx.doi.org/10.1037/0033-295X.105.3.442

Ashby, F. G., & Maddox, W. T. (2005). Human Category Learning. *Annual Review of Psychology*, *56*(1), 149–178. http.//doi.org/10.1146/annurev.psych.56.091103.070217

Haider, H., & Frensch, P. A. (2005). The generation of conscious awareness in an incidental learning situation. *Psychological Research*, *69*(5-6), 399–411. http://doi.org/10.1007/s00426-004-0209-2

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*(2), 106–120. http://doi.org/10.1101/lm.1.2.106

Lagnado, D. A., Newell, B. R., Kahan, S., & Shanks, D. R. (2006). Insight and strategy in multiple-cue learning. *Journal of Experimental Psychology: General*, *135*(2), 162–183. http://doi.org/http://dx.doi.org/10.1037/0096-3445.135.2.162

Price, A. L. (2009). Distinguishing the contributions of implicit and explicit processes to performance of the weather prediction task. *Memory & Cognition*, *37*(2), 210–222. http://doi.org/10.3758/MC.37.2.210

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, *118*(3), 219–235. http://doi.org/http://dx.doi.org/10.1037/0096-3445.118.3.219

Sanchez, D. J., Gobel, E. W., & Reber, P. J. (2010). Performing the unexplainable: Implicit task performance reveals individually reliable sequence learning without explicit knowledge. *Psychonomic Bulletin & Review*, *17*(6), 790–796. http://doi.org/10.3758/PBR.17.6.790

Sergent, C., & Dehaene, S. (2004). Is Consciousness a Gradual Phenomenon? Evidence for an All-or-None Bifurcation During the Attentional Blink. *Psychological Science*, *15*(11), 720–728. http://doi.org/10.1111/j.0956-7976.2004.00748.x

Smith, R. W., & Kounios, J. (1996). Sudden insight: All-or-none processing revealed by speed–accuracy decomposition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(6), 1443–1462. http://doi.org/http://dx.doi.org/10.1037/0278-7393.22.6.1443

# Chapter 1 How effective is incidental learning of the shape of probability distributions?

*Randy Tran, Edward Vul, and Harold Pashler*

ROYAL SOCIETY OPEN SCIENCE

rsos.royalsocietypublishing.org

Research

**Cite this article:** Tran R, Vul E, Pashler H 2017 How effective is incidental learning of the shape of probability distributions? *R. Soc. open sci.* **4**: 170270.
http://dx.doi.org/10.1098/rsos.170270

Received: 29 March 2017
Accepted: 3 July 2017

**Subject Category:**
Psychology and cognitive neuroscience

**Subject Areas:**
cognition/psychology

**Keywords:**
implicit learning, generative models, probability learning

**Author for correspondence:**
Randy Tran
e-mail: r4tran@ucsd.edu

Electronic supplementary material is available online at https://dx.doi.org/10.6084/m9.figshare.c.3832507.

THE ROYAL SOCIETY PUBLISHING

# How effective is incidental learning of the shape of probability distributions?

Randy Tran, Edward Vul and Harold Pashler

Department of Psychology, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA

RT, 0000-0002-2353-6858

The idea that people learn detailed probabilistic generative models of the environments they interact with is intuitively appealing, and has received support from recent studies of implicit knowledge acquired in daily life. The goal of this study was to see whether people efficiently induce a probability distribution based upon incidental exposure to an unknown generative process. Subjects played a 'whack-a-mole' game in which they attempted to click on objects appearing briefly, one at a time on the screen. Horizontal positions of the objects were generated from a bimodal distribution. After 180 plays of the game, subjects were unexpectedly asked to generate another 180 target positions of their own from the same distribution. Their responses did not even show a bimodal distribution, much less an accurate one (Experiment 1). The same was true for a pre-announced test (Experiment 2). On the other hand, a more extreme bimodality with zero density in a middle region did produce some distributional learning (Experiment 3), perhaps reflecting conscious hypothesis testing. We discuss the challenge this poses to the idea of efficient accurate distributional learning.

## 1. Introduction

People often seem to behave effectively based on noisy observations of uncertain environments. This might seem surprising because people generally have poor incidental memory (e.g. the direction that Lincoln faces on the penny; [1]). On the other hand, probability distributions may be special, and there is evidence that people are quite good at estimating frequencies of events even when they have paid little attention to the stimuli as they appeared [2]. One currently popular interpretation of this adaptive flexibility assumes that people efficiently learn probabilistic generative models of their environment and then use these models to guide their behaviour. Such a capability would seem to have the potential to assist people in achieving many of their goals, including goals with strong benefits to Darwinian fitness (such as finding food and finding mates). If one looks at the literature, however, while there are many examples of evidence taken to favour the idea of flexible induction of generative models, the evidence appears somewhat restricted and indirect. For example, Vulkan [3] showed that people were able to match reward probabilities of several alternatives with their choices, indicating that they can learn probability distributions over those alternatives. People can also learn the probabilistic dependency structure in networks of binary variables (e.g. [4,5]). As impressive as these feats are, these outcomes could potentially be achieved by learning only the first few moments (mean and variance) of a distribution rather than the full underlying structure. In this paper, we ask more straightforwardly: can people learn the overall shape of an observed distribution and are they able generate new instances that retain the properties of the learned distribution?

## 2. Prior methodological procedures used in distributional learning studies

To date, studies that have shown evidence for (e.g. [6,7]) and against (e.g. [8]) distributional learning have used tasks that: (i) employ other types of strategies or (ii) seem to allow for an aggregate analysis of only a few moments of the distribution rather than the whole. Our main focus is on how the methodological procedures from various studies might limit the ability to tease apart what properties are learned from a distribution.

Griffiths & Tenenbaum [6] suggested that people have acquired a great deal of information about the shape of the distribution of quantities such as *baking time for cakes, reigns of Pharaohs* and *booking time for telephone ticket booking agencies.* Their argument for this conclusion was based on participants' ability to answer questions of the form 'If you were calling a telephone box office to book tickets and had been on hold for 3 min, what would you predict for the total time you would be on hold?' They found that people's responses generated from their internal generative models were very similar to the true statistical distributions. However, Mozer *et al*. [9] questioned the conclusions of Griffiths & Tenenbaum [6], arguing that the excellent performance at the aggregate level might be consistent with very limited learning at the individual level (cf. [10,11]).

A broader concern with studies like Griffiths & Tenenbaum [6] is that we do not know how much exposure, and what type of exposure, people have had to events like Pharaohs and ticket-service call lines. Many of the quantities considered by Griffiths & Tenenbaum are subject to soft constraints from general world knowledge: knowing current average lifespans have increased over time, knowing that some pharaohs came into power at a very young age all impose constraints on the distribution of Pharaoh reigns. Thus, distributional knowledge about such world facts need not imply that people efficiently learn such distributions from direct observation; instead it may imply that people effectively infer this distribution as needed.

Sailor & Antoine [8] used a more controlled set of stimuli with a task requiring participants to estimate the size of squares drawn from two distributions (Experiments 1, 3 and 4: overlapping; Experiments 2 and 5: non-overlapping). On a given trial, participants were initially presented with a square drawn from one of the distributions and were coloured red or blue to distinguish which distribution the square was drawn from; however, this was never explicitly stated to the participants. Participants then had to estimate the size of the initially presented square by adjusting the size of a subsequently displayed square. Only on the last two trials of the experiment were participants asked to estimate the mean size of the red and blue squares. Sailor & Antoine found that the estimated means for both the red and blue squares did not differ from the average of the two distributional means. In other words, participants were unable to distinguish the two different distributions; instead, they grouped the red and blue squares into a single-unimodal distribution. We argue that the methodology presented by Sailor & Antoine may not well assess an individual's ability to learn the shape of a distribution, because it requires only sensitivity to averages. In a similar task, Gershman & Niv [12] had participants estimate the number of circles presented on the screen. The circles were either all red or all blue and were drawn from two different underlying distributions of quantity. In line with the findings from Sailor & Antoine, Gershman & Niv also discovered that participants' estimations were biased towards the mean of both distributions. Participants' biases, however, were reduced when the red and blue quantity distributions were further apart (i.e. more easily distinguishable; see Experiment 3 of this article for comparable results).

In a further investigation of these findings, Xu & Griffiths [7] were able to show that participants can learn properties of a bimodal distribution using a serial reproduction task. Xu & Griffiths employed a similar procedure to Sailor & Antoine [8] where participants learned to distinguish two types of fish drawn from two separate size distributions. On a given trial, a to-be-estimated fish was presented on the screen and disappeared. Participants then adjusted the size of a subsequent fish to estimate the just-seen fish. A major novelty in Xu & Griffiths' procedure was that each estimation made by a participant was used as subsequent to-be-estimated fish. In other words, participants estimated fish sizes from their own previous estimates (i.e. a Markov chain) rather than estimating fish sizes from fish independently drawn from the experimental distribution on each trial. Using this procedure, Xu & Griffiths claimed to have demonstrated learning of a bimodal distribution. However, the argument rests on people's reconstructions of their own estimates where iterated learning can occur from trial to trial. Hence, with this paradigm, one cannot straightforwardly ask whether or not people can generate new instances that conform to a learned distribution because each trial is influenced by the previous trial.

# 3. Present study

The current study was designed to provide a test as simple and direct as possible for the idea that people implicitly learn the shape of a distribution based on observed samples of that distribution. The study represented something of a (friendly) 'adversarial collaboration' (cf. [13]), in that one of us (EV) was generally favourably disposed to the idea of implicit learning of generative models, while another of us (HP) was fairly sceptical of this idea, and RT at least professed neutrality.

To maximize the chances of demonstrating effective distributional learning, several features were built into the design. First, the variable whose distribution was tested was a variable that was highly relevant to actions the subjects would be performing. To arrange this, we used a 'whack-a-mole'-type game in which the subject sought to click on an object during the brief period before it disappeared. This required paying close attention to its location as the object's sole action-relevant property. Second, we exposed subjects to a distinctive and somewhat unusual (bimodal) distribution to make it possible to test the fidelity of the distribution they learned. The test of learning used here required subjects to produce their own sequence of locations, mimicking the locations observed during the learning phase. While the virtues of this form of test can be debated (see General discussion), the goal here was to maximize the chance of finding distributional learning (see [14], for arguments that the mental representation of distributions is embodied in the ability to generate new samples from these distributions).

# 4. Experiment 1

In Experiment 1, learning was incidental: subjects played the game in Phase 1 with no expectation of being tested.

## 4.1. Method

### 4.1.1. Participants

Thirty undergraduates at the University of California, San Diego participated in this experiment for course credit. All were naive to the purpose of the experiment.

### 4.1.2. Distribution used in phase 1

A single fixed sequence of locations was used for all subjects in Phase 1 (the entire sequence is provided in the electronic supplementary material). The purpose of this was to avoid any confusion of the results due to sampling variability of the observations. The distribution of values used included only multiples of 0.01, with one observed value at each position within the unit interval (0, 1), plus additional values 'piled up' over two modes, one ranging from 0.10 to 0.26 (with four total observations at each point in that range) and the other from 0.80 to 0.84 (with seven total observations at each point in that range). Figure 1 shows this distribution.

### 4.1.3. Procedure in phase 1

In Phase 1, subjects were told 'Welcome to the experiment. In the first phase of the study you will play a game similar to the old computer game *Wackamole*. On every play, a disk will appear and begin expanding. Your job is simply to click on it before it disappears. If you click on it before it disappears, you score. That's it!' They played the game 180 times. Average viewing distance from the screen was about 76 cm on a $1024 \times 768$-resolution screen. On each trial, a blue disc (initially just 1 pixel) appeared in a horizontal range of positions 822 pixels in width centred on the screen. Beginning at the moment

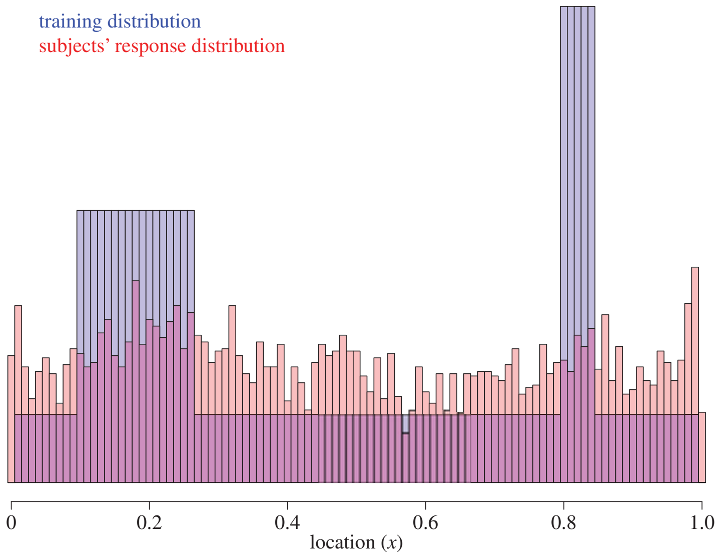


**Figure 1.** Experiment 1 results: histogram of the training (blue) and reported (red) locations.



of its appearance, each disc grew at 100 pixels $s^{-1}$ until it reached a size of 100 pixels, at which time it disappeared. If the subject was able to click on the item during its 1 s expansion phase, they received 1 point. Nothing else appeared on the horizontal line. (If the subject hit the disc, a confirming sound would play with the word 'Hit' displayed on the screen. Otherwise a buzz sound played while the word 'Miss' was displayed. The feedback lasted 1.8 s. A streak counter and best streak counter were also visible on the top left of the screen, displaying the subject's current hit streak and their best hit streak overall.)

#### 4.1.4. Procedure in phase 2

Immediately after the last play, subjects began the second phase of the study, and were told, 'Now we are interested in determining how much of an intuitive sense you have gained for how the locations of the disks were being determined. Please show us this by generating a new **sequence of locations**. Please do NOT click in the same spot over and over.' They were also told 'If you think there were any other patterns in the original sequence, please try to generate a sequence that reflects those patterns, too. **Don't worry about mimicking the timing** of the original sequence. Just try to produce a **sequence of locations** which is as much like the original sequence as you can make it.'

In the second phase, subjects' clicks were self-paced. When they clicked, a disc showed up with the location of the cursor as the centre of the disc, a click counter on the top left of the screen incremented with each click. After 180 clicks were registered, an exit screen was displayed, terminating the study.

### 4.2. Results and discussion

The average hit rate of clicks in Phase 1 was $M = 0.51$, s.d. $= 0.14$, s.e. $= 0.02$. Scores ranged from 0.25 to 0.74 with a median of 0.50. As with the real Whack-a-mole game, we expected to find a wide range of hit rates during Phase 1. Figure 1 shows the distribution of generated click positions aggregated across subjects for Phase 2. The subjects' responses show no obvious similarity to the bimodal pattern presented in Phase 1.

## 5. Experiment 2

In Experiment 2, the task was the same, but the subjects were warned that they would be tested on the distributions of locations.

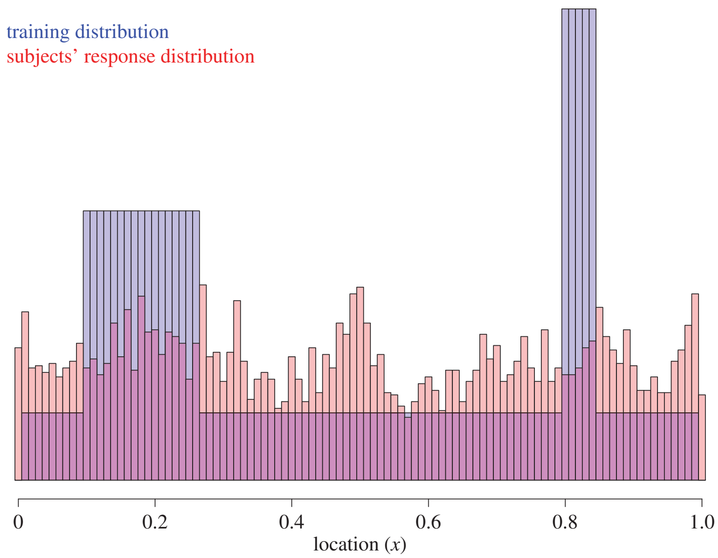


**Figure 2.** Experiment 2 results: histogram of the training (blue) and reported (red) locations.

## 5.1. Method

### 5.1.1. Participants

Thirty-one undergraduates drawn from the same population as Experiment 1 participated. All were naive to the purpose of the experiment.

### 5.1.2. Materials and design

Materials and design were identical to Experiment 1 with the exception of a difference in instructions.

### 5.1.3. Procedure

The procedure was identical to that of Experiment 1 except that prior to performing the first phase (playing Whack-a-mole), the subjects were told: 'Just one more thing: please pay attention to the sequence of locations where the disk appears. After you're done playing, we'll ask you to try to generate a sequence of locations that simulates the sequence the computer is generating. So please see if you can learn the characteristics of the sequence of locations where the disks pop up.'

## 5.2. Results and discussion

One subject was excluded from the subsequent analyses due to a logging error in the subject's file. The average hit rate of clicks in Phase 1 was $M = 0.57$, s.d. $= 0.18$, s.e. $= 0.03$. Scores ranged from 0.17 to 0.89 with a median of 0.60. Figure 2 shows the subjects' response distribution for Phase 2. Again, there was no sign in the aggregate responses that subjects learned the bimodality of the distribution, despite an explicit instruction to try to learn the characteristics of the sequence.

# 6. Experiment 3

In Experiment 3, the procedure followed Experiment 1, except the distribution used was more extremely bimodal, with zero density outside of the intervals of the modes ([0.10, 0.26] and [0.8, 0.84]; figure 3). Thirty-one undergraduates from the same subject pool participated (one subject was excluded due to a file logging error). The complete stimulus sequence is provided in the electronic supplementary material.

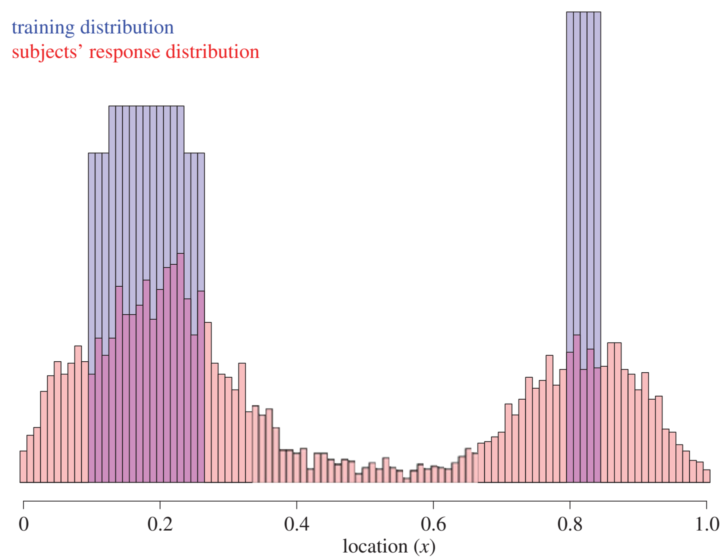


**Figure 3.** Experiment 3 results: histogram of the training (blue) and reported (red) locations.

### 6.1. Results and discussion

The average hit rate of clicks in Phase 1 was $M = 0.58$, s.d. $= 0.16$, s.e. $= 0.03$. Scores ranged from 0.24 to 0.83 with a median of 0.63. Figure 3 shows that here in Phase 2, subjects did indeed pick up on the now discrete bimodality much better than in the previous studies.

## 7. Quantifying learning

While one of us (HP) felt that the results clearly showed that distributional shape learning was negligible except when the distribution had a gross qualitative feature (zero density in the middle region), EV felt it would still be useful to explore the extent of learning quantitatively.

To characterize learning in these experiments, we therefore ascertained which precision of a kernel density estimate applied to the training observations ($x_i$ being the position seen on a given training trial) best captured the responses produced by our observers. As the range of possible responses is bounded, we created a 'Beta kernel' parametrized by one precision parameter, $k$:

$$\hat{f}_k(x) = \frac{1}{n}\sum_{i=1}^{n} K_k(x|x_i)$$

and

$$K_k(x|x_i) = \mathrm{Beta}(x|1 + x_i \times 10^k, 1 + (1 - x_i) \times 10^k)$$

This fitting was done by obtaining the distribution over positions as the kernel density estimate (sum over all kernels for all training data, normalized), for a given $k$. Then the likelihood of a subject's responses under that distribution was calculated for each $k$. Finally, the maximum-likelihood $k$ was taken as the estimate. When $k$ is large (greater than 0), the kernel amounts to a beta distribution peaked at the observed value, with the distribution approaching a single spike at the observed value as $k$ increases. When $k$ is small (less than 0), the kernel loses just about all of the information about the observed value, and yields a uniform (0, 1) distribution (figure 4). We fit the kernel precisions parameter to individual subjects in each of our experiments, as well as the aggregate across-subject data shown in figures 1–3.

Figure 5 shows how well different values of $k$ fit individual subjects in each of our experiments. Only Experiment 3 shows that subjects learned something from the training distribution—as indicated by an advantage of kernel precisions greater than 0 (23 of 30 subjects have a best-fitting $k_{\mathrm{ML}} > 0$). By contrast, for Experiments 1 and 2, the best-fitting kernel precision is very negative for most subjects ($k_{\mathrm{ML}} > 0$ for 10/30 and 9/30, respectively), indicating that most subjects' responses reflect effectively zero influence of

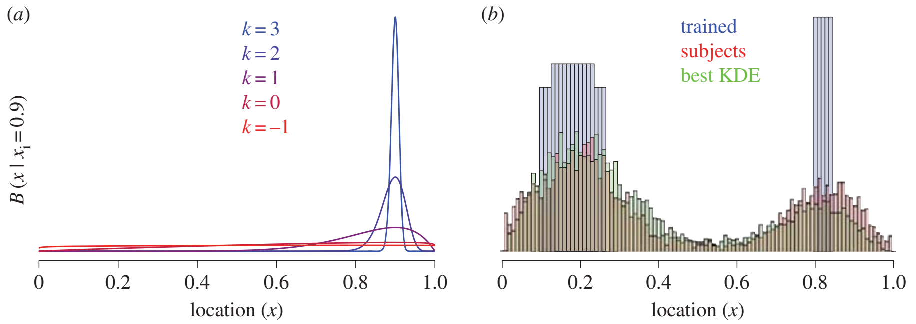


**Figure 4.** (*a*) Beta kernel for an observed value of 0.9 for different precision parameters. Kernel precision parameters less than 0 yield an effectively uniform distribution. (*b*) Results of Experiment 3: trained distribution (blue), subjects' response distribution (red), and the best-fitting ($k = 1.09$) Beta kernel density estimate of the trained distribution (green; note yellow arises where red and green overlap).

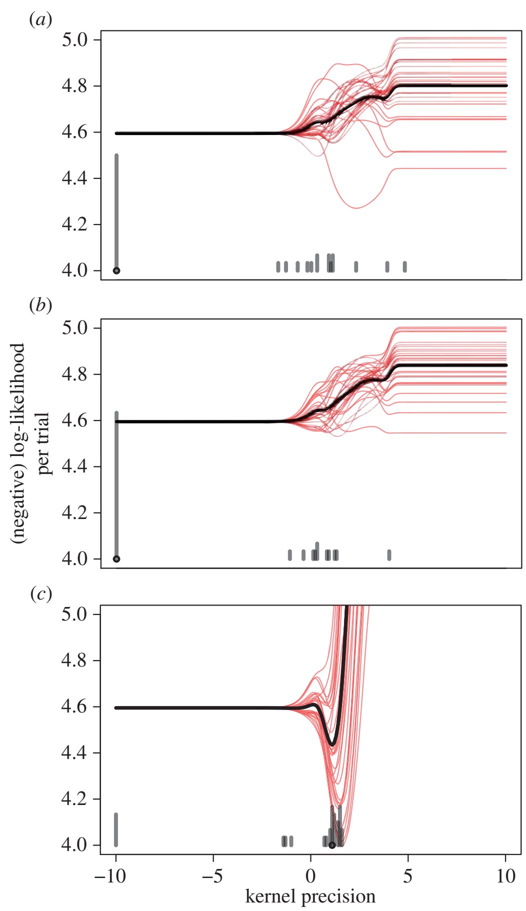


**Figure 5.** Quality of fit (*y*-axis; lower is better) of different kernel precision values (*x*-axis) for Experiments 1, 2 and 3 (from (*a*) to (*c*)). Individual subject fits are shown in red, while the fit to the aggregate data is shown in black. Grey bars at the bottom of each panel are a histogram (across subjects) of the best-fitting *k*-values (black circle indicates the best-fitting value for the aggregate over all subjects). Although a minority of subjects reveal some learning (positive kernel precision) in Experiments 1 and 2, for the most part, kernel precisions are very negative, indicating that subjects do not reliably capture any of the training distribution signal in their responses. By contrast, Experiment 3 shows reliable learning.

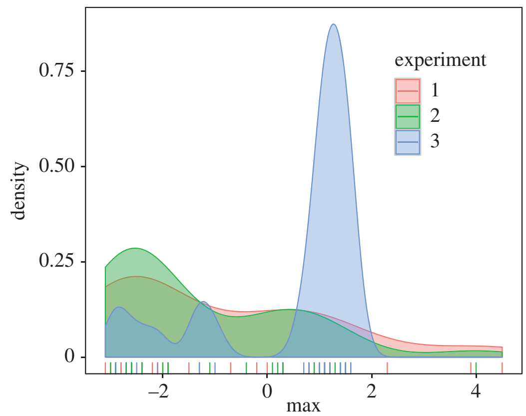


**Figure 6.** Histogram of subject MAP estimates for Experiments 1–3.

the training distribution. To examine these results across experiments, we estimated maximum *a posteriori* values of $k$ (with loosely informative priors of $k \sim N(0, 5)$ to avoid indistinguishable regions for very negative values of $k$) for each subject. A one-way ANOVA on subject MAP estimates (figure 6) showed statistical significance across the three experiments, $F_{2,89} = 8.54$, $p < 0.001$.

## 8. General discussion

The results reported here show that when people are exposed to a stream of stimuli whose properties on an action-relevant dimension (here, location) conform to a bimodal distribution, they fail to spontaneously learn the bimodality. The same is true even when they are told to try to learn the distribution of locations (Experiment 2). However, when the continuous bimodal distribution was discretized by adding a zero-density gap between the two modes, people showed clear distributional knowledge (Experiment 3). We are inclined to reject the possibility that subjects learned the distribution in Experiments 1 and 2, but failed to produce under the 'generate samples' testing procedure because the same procedure yielded clear positive findings in Experiment 3.

The results are amenable to a number of possible interpretations, some of which we will mention here without taking any strong view (reflecting the current authors' friendly adversarial collaboration noted in the Introduction). One intriguing interpretation is that there is no general non-parametric learning of continuous probability distributions, and the results of Experiment 3 arose because the distribution could be readily discretized on account of the zero-density interval between the two modes. This would be consistent with the idea that discreteness affects the ease in which distributions can be abstracted.

Another possibility is that people have a strong tendency to learn by 'parameter tuning' of particular functional forms of distributions, rather than learning distributions non-parametrically—an account echoed in results suggesting that perceptual learning amounts to parameter tuning of feature relationships, rather than learning new relationships among features [15]. A more mundane but still interesting possibility is that non-parametric learning of a distribution proceeds with imperfect and incomplete memory, which renders the distributions in Experiments 1 and 2 too subtle to be learned. Although learning does occur when the subtle bimodal distribution was made more notable in Experiment 3. These findings contrast the conclusions from Acerbi *et al*. [16] where the discernibility of complex distributions does not modulate performance. However, their subjects were given explicit distribution information to be used in their spatial estimation task. This discrepancy may reflect different processes when using explicit information versus generating novel samples from distributions. Yet another possibility, attributed to a referee suggestion on an earlier version of this manuscript, might be subjects represent the distribution faithfully as changes in distance from trial-by-trial (i.e. allocentrically) rather than over the spatial width of the line. Finally, a related possibility is that people have prior assumptions strongly favouring unimodal distributions, and the data provided in Experiments 1 and 2 (perhaps corrupted by memory) are insufficient to overcome such priors.

Further support for this possibility comes from a recent study by Sanborn & Beierholm [17]. These investigators had subjects estimate the number of circles in a display (the number ranged from a minimum of 23 to at most 35). Quantities were drawn from discretized bimodal or quadrimodal distributions and feedback was provided. While Sanborn & Beierholm's bimodal distribution did not have zero density in the middle, by contrast with our Experiment 3, the two modes were always at the most extreme left and right positions within the distribution (e.g. trials with 23 and 29 circles appeared with a probability of 0.3 and trials with 24–28 circles appeared with a probability of 0.08 with a total of 700 trials). The quadrimodal distribution combined two bimodal distributions with a zero density region interposed between them (e.g. trials with 23, 25, 29 or 31 circles appeared with a probability of 0.2; trials with 24 or 30 circles appeared with a probability of 0.1; all other possible quantities of circles had a probability of 0). Subjects' behavioural responses (shown as conditional response distributions) suggested they had learned a good deal about the distributions. Given their discrete character, these findings seem consistent with the findings of Experiment 3.

## 9. Suggestions for future research

So what do these three studies tell us? It seems that learning fine-grained structure of observed probability distribution may not be as efficient as prior literature might seem to imply. The clear discrepancy between Experiments 1 and 2, and Experiment 3 suggests an intriguing possibility: only when a continuous distribution may be easily discretized do people engage in some form of non-parametric learning; otherwise, they tend to learn only a few moments of the distribution (such as the commonly investigated tendency to learn the mean and variance).

Of course, before such a hypothesis might be acknowledged, it would be important to untangle more mundane accounts: perhaps strong priors about unimodality, coupled with an imperfect memory for exemplars, is responsible for this pattern of results.

Ethics. The research was approved by the University of California San Diego Social and Behavioral Sciences Institutional Review Board. All subjects provided written informed consent for their participation.

Data accessibility. Our data are deposited at FigShare: https://doi.org/10.6084/m9.figshare.4767997.

Authors' contributions. H.P. principally conceptualized and designed the study. R.T., E.V. and H.P. contributed to analysis and interpretation of the data, contributed to drafting and revising the manuscript. All authors gave final approval for publication.

Competing interests. We declare we have no competing interests.

Funding. This work was supported by a collaborative activity award to H. Pashler from the J. S. McDonnell Foundation, a MURI award from the Office of Naval Research (25684A), and two NSF Grants (SBE-0542013, G. W. Cottrell, PI; SES-1461535, M. Mozer, PI).

Acknowledgements. We thank Kevin Ngo for programming assistance and Noriko Coburn for overseeing data collection.

## References

1. Nickerson RS, Adams MJ. 1979 Long-term memory for a common object. *Cognit. Psychol.* **11**, 287–307. (doi:10.1016/0010-0285(79)90013-6)
2. Hasher L, Zacks RT, Rose KC, Sanft H. 1987 Truly incidental encoding of frequency information. *Am. J. Psychol.* **100**, 69–91. (doi:10.2307/1422643)
3. Vulkan N. 2000 An economist's perspective on probability matching. *J. Econ. Surv.* **14**, 101–118. (doi:10.1111/1467-6419.00106)
4. Cheng PW. 1997 From covariation to causation: a causal power theory. *Psychol. Rev.* **104**, 367–405. (doi:10.1037/0033-295X.104.2.367)
5. Griffiths TL, Tenenbaum JB. 2005 Structure and strength in causal induction. *Cognit. Psychol.* **51**, 334–384. (doi:10.1016/j.cogpsych.2005.05.004)
6. Griffiths TL, Tenenbaum JB. 2006 Optimal predictions in everyday cognition. *Psychol. Sci.* **17**, 767–773. (doi:10.1111/j.1467-9280.2006.01780.x)
7. Xu J, Griffiths TL. 2010 A rational analysis of the effects of memory biases on serial reproduction. *Cognit. Psychol.* **60**, 107–126. (doi:10.1016/j.cogpsych.2009.09.002)
8. Sailor KM, Antoine M. 2005 Is memory for stimulus magnitude Bayesian? *Mem. Cognit.* **33**, 840–851. (doi:10.3758/BF03193079)
9. Mozer MC, Pashler H, Homaei H. 2008 Optimal predictions in everyday cognition: the wisdom of individuals or crowds? *Cogn. Sci.* **32**, 1133–1147. (doi:10.1080/03640210802353016)
10. Lewandowsky S, Griffiths TL, Kalish ML. 2009 The wisdom of individuals: exploring people's knowledge about everyday events using iterated learning. *Cogn. Sci.* **33**, 969–998. (doi:10.1111/j.1551-6709.2009.01045.x)
11. Vul E, Goodman N, Griffiths TL, Tenenbaum JB. 2014 One and done? Optimal decisions from very few samples. *Cogn. Sci.* **1**, 107–143. (doi:10.1111/cogs.12101)
12. Gershman SJ, Niv Y. 2013 Perceptual estimation obeys Occam's razor. *Front. Psychol.* **4**, 623. (doi:10.3389/fpsyg.2013.00623)
13. Mellers B, Hertwig R, Kahneman D. 2001 Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychol. Sci.* **12**, 269–275. (doi:10.1111/1467-9280.00350)
14. Griffiths TL, Vul E, Sanborn AN. 2012 Bridging levels of analysis for probabilistic models of cognition. *Curr. Dir. Psychol. Sci.* **21**, 263–268. (doi:10.1177/0963721412447619)
15. Michel MM, Jacobs RA. 2007 Parameter learning but not structure learning: a Bayesian network model of constraints on early perceptual learning. *J. Vis.* **7**, 4. (doi:10.1167/7.1.4)
16. Acerbi L, Vijayakumar S, Wolpert DM. 2014 On the origins of suboptimality in human probabilistic inference. *PLoS Comput. Biol.* **10**, e1003661. (doi:10.1371/journal.pcbi.1003661)
17. Sanborn AN, Beierholm UR. 2016 Fast and accurate learning when making discrete numerical estimates. *PLoS Comput. Biol.* **12**, e1004859. (doi:10.1371/journal.pcbi.1004859)

Chapter 1, in full, is a reprint as it appears in How effective is incidental learning of the shape of probability distributions? Tran, Vul, and Pashler (2017). The dissertation author was the primary investigatory and author of this paper.

# Chapter 2 **Learning to exploit a hidden predictor in skill acquisition: Tight linkage to conscious awareness**

*Randy Tran and Harold Pashler*

PLOS | ONE

RESEARCH ARTICLE

# Learning to exploit a hidden predictor in skill acquisition: Tight linkage to conscious awareness

**Randy Tran*, Harold Pashler**

Department of Psychology, University of California, San Diego, San Diego, California, United States of America

* r4tran@ucsd.edu

Check for updates

OPEN ACCESS

**Citation:** Tran R, Pashler H (2017) Learning to exploit a hidden predictor in skill acquisition: Tight linkage to conscious awareness. PLoS ONE 12(6): e0179386. https://doi.org/10.1371/journal.pone.0179386

**Editor:** J Bruce Morton, Western University, CANADA

**Received:** February 16, 2017

**Accepted:** May 28, 2017

**Published:** June 20, 2017



**Data Availability Statement:** All data are available from the FigShare database: https://doi.org/10.6084/m9.figshare.4654783.

**Funding:** This work was supported by the Office of Naval Research (MURI Grant #N00014-10-1-0072), a collaborative activity grant from the James S. McDonnell Foundation, and the National Science Foundation (Grant SBE-582 0542013 to the UCSD Temporal Dynamics of Learning Center and Grant SES 1461535 to PI: Michael Mozer). The funders had no role in study design, data collection and

## Abstract

It is often assumed that implicit learning of skills based on predictive relationships proceeds independently of awareness. To test this idea, four groups of subjects played a game in which a fast-moving "demon" made a brief appearance at the bottom of the computer screen, then disappeared behind a V-shaped occluder, and finally re-appeared briefly on either the upper-left or upper-right quadrant of the screen. Points were scored by clicking on the demon during the final reappearance phase. Demons differed in several visible characteristics including color, horn height and eye size. For some subjects, horn height perfectly predicted which side the demon would reappear on. For subjects not told the rule, the subset who demonstrated at the end of the experiment that they had spontaneously discovered the rule showed strong evidence of exploiting it by anticipating the demon's arrival and laying in wait for it. Those who could not verbalize the rule performed no better than a control group for whom the demons moved unpredictably. The implications of this tight linkage between conscious awareness and implicit skill learning are discussed.

## Introduction

A critically important aspect of skill acquisition is learning to take advantage of the various predictive relationships that exist within the relevant domain. Through reinforcement learning and other learning processes, people are assumed to discover and exploit such predictive relationships and to optimize their performance accordingly, achieving greater rewards as their skill grows (e.g., [1]). The question posed in the present article is: do people learn to exploit predictive relationships without showing any conscious awareness of the relationship that they are exploiting?

Anyone acquainted with the cognitive psychological literature might suppose that the answer to this question is clearly "yes". Indeed, there are several well-known lines of research which seem to show beyond any doubt that implicit procedural learning takes place without conscious awareness (see [2], for a review; c.f., [3]). Moreover, in the past 15 years or so, rather little in the way of new research on the topic seems to have been published, possibly suggesting

analysis, decision to publish, or preparation of the manuscript.

**Competing interests:** The authors have declared that no competing interests exist.

that the issue has seen by many as "settled". However, as will be seen below, these results, though intriguing, involve very limited kinds of behavioral changes that are not necessarily representative of skill learning in the broader sense. The remainder of this introduction provides a brief overview of research demonstrating unconscious procedural learning, pointing out how these studies leave open the general question posed above. We then go on to describe the construction of a very simple videogame designed expressly to revisit and shed light on the question posed here.

## Evidence for unconscious implicit learning

A number of experimental designs have produced results that appear to show extensive implicit learning without awareness. In the best known of these studies, Willingham, Nissen, and Bullemer [4] had people make a series of button pushes in response to a spatial sequence of stimuli, and repeated a 10-item sequence throughout the experiment. Subjects were divided into 3 different groups depending on their ability to explicitly describe the repetition. Most subjects showed a speed-up in performance for repeated sequences, but were unable to verbally describe the sequences that repeated. While the reduction in RT with practice was greater for subjects who could describe the repeated sequence, there was substantial improvement even for those who could not describe it (for discussion, see [5–9].)

Another commonly cited line of evidence for learning of predictive relationships without awareness comes from Miller [10,11], who had subjects respond to the identity of a central "target" characters while ignoring some other "flanker" characters that were presented on either side of the central character. The identity of the flanker characters partially predicted the identity of the target character. Subjects responded faster on trials that conformed to this predictive relationship, as compared to trials that deviated from it ([10,11]). Several pieces of evidence argued for a dissociation between awareness of the predictive relationship and behavioral reliance upon the relationship. For example, a small proportion of subjects were unable to report the most common flanker-character pairing, and for these "unaware" subjects the flanker effect on response latencies was actually *stronger*, rather than weaker, than for other subjects (although the difference was not statistically significant.)

While these results would appear to suggest that unconscious implicit learning is probably ubiquitous, the research designs represent only a rather narrow set of behavioral changes compared to the typical real-world skill acquisition challenge that people face. For one thing, the choice reaction-time tasks used in some studies required people to respond quickly and accurately to the stimuli that were only partially predictable based on the covert relationship (in the Willingham et al. [8] design, that was based on the identity of the previous stimulus; in the Miller [10,11] studies, that was the identity of the flanker.) Exploiting the presence of these stimuli may therefore have involved tuning the perceptual system itself to lower the threshold for identifying stimuli likely given the context. This seems quite different than choosing overt actions taken based on anticipations of predictable future events. It is also possible that people might have failed the tests of awareness given by these investigators because these tests required them to produce information that was valid only on some proportion of trials (subjects might not comment on a regularity they had consciously noticed at one point in learning if they had later observed apparent disconfirmation of the regularity).

Moreover, the repetition of the stimulus series in the Willingham, Nissen, and Bullemer [8] design also meant that a sequence of responses was repeated. Subjects may have formed higher level "chunks" of motor programs that represented multiple finger responses. These may be poorly verbalized precisely because they represent motor response patterns (just as we would not expect people to be well able to verbalize what they do when they play ping-pong.)

## Empirical challenges to unconscious learning

Another set of studies from Lewicki and colleagues would appear to be more directly on point for the question posed above—and relatively immune to the objections raised in the preceding paragraph. Lewicki and colleagues published numerous studies that purported evidence for learning hidden covariations (e.g., [12]; see also [13] for a review.). For example, in Hill et al. [12], subjects were shown faces that covaried facial features and personality characteristics (e.g., "fair professors always had 'long' faces. . .unfair professors always had 'short faces'.") in the training phase. Next, subjects rated the "fairness" of novel faces in the testing phase. The data showed that subjects tended to rate novel long faces to be fair and novel short faces to be unfair (Hill et al. [12], Experiment 1). However, the subjects were unable to verbalize the hidden covariation in their exit survey and were therefore said to have acquired it unconsciously. The evidence from Hill et al. suggest unconscious acquisition of hidden covariation can be exploited for future events.

Unfortunately, however, there is reason to doubt the replicability of these studies. Hendrickx, de Houwer, Baeyens, Elen, and Van Avermaet [14] attempted 9 conceptual and 3 direct replications and only one of these efforts confirmed the original finding. Hendrickx et al. [14] suggested that their replication attempts were actually better controlled than the original studies (i.e., minimized correlated features) and had more sensitive awareness measures (e.g., a recognition questionnaire with elaboration instead of free response). Hendrickx et al.'s [14] replication of the described study showed that only subjects who were able to describe the hidden covariation in the exit survey showed the predicted pattern of results (e.g., long faces rated as fair, short faces rated as unfair). Subjects who were unable to describe the covariation did not show this effect. Given the large number of non-replications, we would suggest that the results of Lewicki and colleagues should not be assumed valid. Moreover, the Hill et al. [12] designs also provided only very coarse measures of awareness; in the studies described here, we provide a more fine grained temporal window on this process.

## Current approach

As described above, most past studies of implicit learning of predictive regularities have focused on very austere choice-reaction time designs in which subjects' task is merely to make one of a small number of button-press responses chosen in compliance with a fixed stimulus-response mapping provided by the experimenter, or on tasks in which people make repeated sequences of motor responses. We felt that by moving to a computer game environment we could examine the effect of embedding hidden predictive relationships on a more robust and compelling example of skill learning.

The very simple videogame used in both experiments described below was designed with two specific aims in mind. One aim was to introduce a predictive regularity that would be valid on 100% of trials, to reduce the likelihood that people would abandon whatever conscious access they might achieve due to encountering contrary cases. Second, rather than relying on subtle latency changes due to priming as in the studies described above, the game was devised so that detection of a regularity would allow the participant to improve his or her score by making active behavioral choices in anticipation of what would happen next in the game environment.

## Experiment 1

The simple computer game used in Experiment 1 worked as follows (Fig 1). Each play began with a simple cartoonlike figure (a "demon") moving upward from the bottom of the screen to the location shown in Fig 1 (point A). At this point, the demon came to a rest, and remained

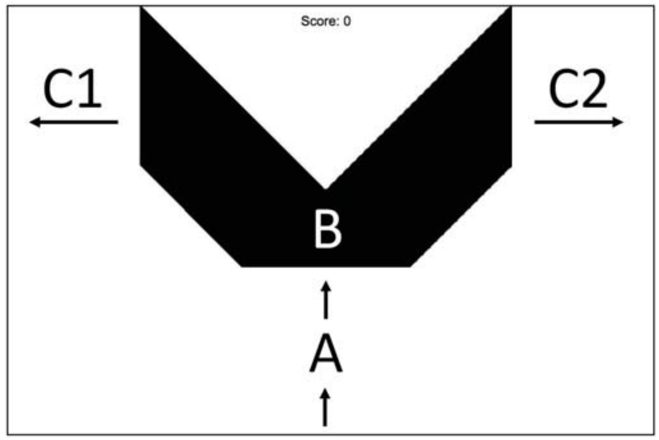


**Fig 1. Example of the game screen.** General appearance of the screen viewed by subjects in the "Demon Hunting Game". The "demon" is paused at the bottom of the screen (A), waiting for the subject to click on it. When the subject does this, the demon resumes its upward trajectory which takes it behind the black V-shaped occluder (B), whereupon it finally emerges moving more rapidly in the location labeled C1 or C2. At this point, the subject must click on the demon before it leaves the screen in order to score. Neither the lines nor the letters are present on the game screen. Unbeknownst to some groups of subjects, the height of the demon's horns predicts whether the demon reappears on the left or right side.

https://doi.org/10.1371/journal.pone.0179386.g001

there until the subject placed the computer mouse cursor over the demon and clicked on it. The purpose of this requirement was to insure that subjects fixated briefly on the demon. Once the mouse click was registered, the demon resumed its upward motion, moving behind a large black V-shaped occluder (point B in Fig 1). After 3 seconds, it re-emerged, either on the left side of the screen heading leftward (point C1 in Fig 1), or on the right side of the screen heading rightward (point C1 in Fig 1). In this phase, it was now moving quite fast (217 pixels/sec), and the player's challenge was to click on the demon before it left the screen. If they could do so, they would earn one point. Subjects performed up to 180 separate "plays", each lasting about 10.5 seconds. Between blocks of 30 plays, subjects were required to rest for 20 seconds without opening any other browser windows (their scores on all blocks of play completed thus far were displayed during this rest period.)

Pilot work in our lab looking in detail at behaviors emerging in this game showed that subjects who were told the rule found it easy to reach near-perfect scoring levels on every trial, and that they accomplished this (as expected) by moving the cursor in advance to the place where the demon was to reappear, and laying in wait for it. Unlike the graded latency changes seen in the implicit learning literature described above, this strategy did not produce a mere subtle modulation of response latency, but rather a change in strategy resulting in a drastic improvement in the rate of scoring (raising performance from not much better than 50% up to nearly 100% in most cases).

In Experiment 1, subjects were randomly assigned to one of four groups. For the first group (Control), whether the demon went left or right was random and not predictable based on any property of the demon. Thus for the Control Group, reliably anticipating the location of reappearance was impossible. For Groups 2–4, the height of the demon's horns perfectly predicted where the demon would reappear: long horns meant that it would go left and short horns meant that it would go right. The difference in horn height was a very salient 5:1. For Group 2 (Predictable/No Instruction Group), the instructions did not mention anything about the predictive relationship. The final two groups of subjects were told either simply that horn height would be relevant (Dimension Instruction Group) or they were given a precise description of the exact rule (Full Instruction Group).

The experiment was divided into 6 blocks of 30 plays. All subjects were told that if they were able to score a point on *every single play* within a given block of 30, their participation would be complete at the end of the block, and they would be paid as soon as they answered a few final questions ("exit interview"). In this exit interview, subjects in all groups (except the Full Instruction group) who reached the performance threshold were asked if they had any hunches enabling them to predict which way the demon would go. The opportunity to be excused from the study after a perfect-scoring block served two purposes: it motivated the subjects to do as well as they could, and it insured that exit interviews took place only a short time after the moment at which the subject first demonstrated mastery of the game.

## Method

### Participants

Subjects were drawn from our laboratory's on-line research subject pool, which provides a diverse panel of subjects of various ages from a wide variety of countries. The subjects provided written informed consent, and the research was approved by the University of California San Diego Social and Behavioral Sciences Institutional Review Board. Subjects are pre-screened for comprehension of English, careful attention to instructions and conscientious performance in prior experiments. A total of 97 subjects completed the study in return for payment of $6.00. Subjects were randomly assigned to one of the four between-subject conditions.

### Stimuli

The experiment was created using Flash web programming IDE; the program ran on client machines and intermittently sent data back to the lab webserver using the JSON protocol (source code is available on request). Demons differed on three dimensions: eye diameter, color of body, and horn height (see Fig 2 for an example). Eye diameter was a random number from a uniform distribution in the range (5 pixels, 35 pixels). Demon bodies were randomly assigned a color from the set {red, green}. For Conditions 2, 3, and 4, horn heights were 10 pixels for demons that went right and 50 pixels for demons that went left. For Condition 1, horn height was chosen from the same set, but it did not predict anything about the demon's behavior. In all cases, horn width was a constant 20 pixels.

**Fig 2. Example stimuli.** Example of a demon used in the game.

https://doi.org/10.1371/journal.pone.0179386.g002

## Procedure

Subjects began by reading instructions. All subjects were given a multiple choice quiz on each element of the instructions that they had read. Subjects who made any mistakes on this quiz were required to reread the instructions and retake the quiz, a process that was repeated until they responded perfectly. For subjects in Condition 1 and 2 (Control and Predictable/Unin-structed groups) the instructions described the goal of the game but said nothing about the location of reappearance of the demon. For subjects in Condition 3 (Dimension Instruction), the instructions stated "One important thing you should know: the height of the demon's horns predicts something about the demon's behavior." For subjects in Condition 4 (Full Instructions), the instructions stated "One important thing you should know: the height of the demon's horns predicts whether the demon will re-appear on the left or right. Long-horned demons reappear on the left, and short-horned demons reappear on the right." (Each of these elements was included on the comprehension quiz given to this group of subjects.)

Subjects then played for 180 trials (unless they achieved a perfect score on any block, at which point their play was terminated). Finally, all subjects (except those in the Full Instruction condition) saw a screen asking them to list up to three hunches they might have about how they could predict which direction the demon would go. For each response, the subject was asked to indicate a level of confidence (using increments of 10%.)

## Results and discussion

To be included in the analysis, subjects had to finish the experiment and answer the survey questions. Data from two subjects (one from the Full Instruction Condition, one from the Dimension Instruction Condition) who completed the experiment were dropped from further analysis because it turned out that they were missing data from more than four trials in one of the blocks due to internet connectivity problems. A few other subjects had missing trials for the same reason, but were missing fewer than four trials total. Those data were included in the analysis. There was a significant main effect of group, $F(4, 97) = 21.76$, $p < 0.001$, and block, $F(5, 485) = 38.35$, $p < 0.001$, as well as a significant interaction, $F(20, 485) = 3.10$, $p < 0.001$.

## Control group performance

The Control group, for whom prediction was not possible, contained 21 subjects. In Fig 3, the line labeled Control Group shows the average performance of this group of subjects (as well as the standard error). The average performance shows no more than a very gradual rise in score level. Thus was expected, because when a subject has failed to anticipate the position of the demon (as would have been the case in at least 50% of plays), it is only occasionally possible to get the mouse over to the appropriate side of the screen rapidly enough to click on it before it disappears—and continued practice produces for most subjects only a limited improvement in score levels achieved by this strategy. A small fraction of subjects do manage this, however, and indeed, two of the 21 subjects (10%) in this condition reached the criterion of perfect performance on a block (neither reported any hunches about predicting the demon's trajectory). The Supplementary Online Materials (S1 File) show the verbatim exit-interview responses of the 17 subjects in this group who were willing to comment in response to the request to specify any rule they thought might possibly have predicted the direction the demon would go in (of course, the direction was actually chosen randomly.) The average confidence reported by the 17 subjects who provided hunches was 59%. Interestingly, there were many highly confident reports of predictive rules that had no valid correspondence to the rules that generated the stimuli (e.g., 70% confidence in "The bigger ears came from the right, whereas the smaller ears ones came from the left." and 70% confidence in "the more i miss on one side is the more it

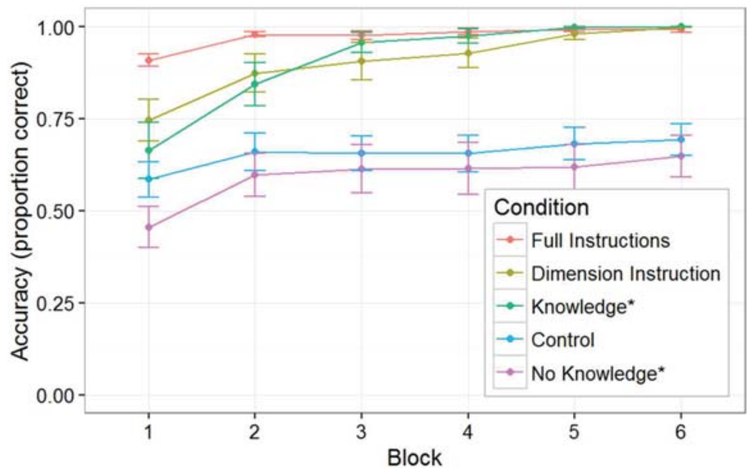


**Fig 3. Experiment 1 subject performance.** Proportion of plays resulting in successful scores for Control, Dimension Instruction, and Full Instruction Groups, as well as for two sub-groups of the Predictive/No-Instruction groups (Knowledge* subjects who were able to verbally report rule, and No-Knowledge* subjects who could not do so.) As described in text, for computing these averages, whenever subjects scored on 100% on a 30-trial block, they were assigned 100% score for computing average performance in the subsequent blocks. No-Knowledge* subjects scores are not superior to those of control subjects for whom there was no predictive relationship available to be exploited.

https://doi.org/10.1371/journal.pone.0179386.g003

goes to that side" and 90% confidence in "coming on to the end of each trial the demons would go either to the left or right about 3 to 5 times straight".)

## Full instructions/dimension instructions group performance

As expected, subjects who were explicitly told the predictive rule or relevant dimension had near perfect mastery by the end of the game. Recall that this game is not a 2-alternative forced choice task, thus some variability can be introduced through motor errors or trials where responses were withheld by the subject.

## Predictable/No instructions group performance

The group whose performance is of greatest interest for examining behavior/awareness dissociations is the Predictable/No Instruction group, containing 26 subjects. Out of these, 14 (54%) were able to perform perfectly on a block of trials, allowing them to terminate their participation. The Supplementary Online Materials show the verbatim exit interview responses of all subjects in the Predictable/No-Instruction group, classified (blind to other aspects of the subject's behavior) according to whether the response indicated complete and accurate knowledge of the rule or not. Two subjects' hunch texts were judged unclassifiable because they contained partial bits of correct information combined with elements of misinformation, thus were excluded from analysis.

In Fig 3, the line labeled No-Knowledge* shows the performance of the 11 subjects judged to have shown no conscious knowledge of the rule. As can be seen, there is no indication that they perform any better than subjects in the Control Group for whom there was no predictive rule, $F(1, 33) = 0.73$, p = 0.40. If they learned anything about the predictive relationship present in the game, they evidently made essentially no use of that learning. Of these 11, one achieved perfect performance in a block, a rate (9%) very similar to the proportion of control subjects achieving this level.

In the Predictable/No-Instructions group, 13 subjects did succeed in describing the rule. All 13 of these Knowledge* subjects (100%) attained perfect performance in one of the blocks of trials. Their performance is shown with the line labeled Knowledge* in Fig 3. Note that for purposes of this graph, whenever subjects scored on 100% of a 30-trial block, they were assigned 100% score for all subsequent blocks—otherwise the rightmost points on graph would reflect an increasingly truncated sample as subjects are peeled off due to having reached perfect performance in an earlier block. (This decision seemed sensible because pilot experimentation in which subjects were required to complete all 6 blocks regardless of performance showed that after people once attained mastery in a block, they scored on close to 100% of trials thereafter; we suspect that the few failures reflected inattention caused by the boredom of performing a task that now lacked any challenge.)

## Individual knowledge* subjects' performance

To provide a more fine-grained picture of the Knowledge* subjects' performance, Fig 4 displays the scores of all 13 individuals in this group block by block. Subjects are highly variable in their Block 1 performance. The figure also shows what appears to be relatively steady progress over the course of the session by most of the subjects who ultimately attain perfect mastery. The latter two statements are jointly confirmed by the strong negative correlation seen between (a) scores on Block 1 and (b) the block number in which the subject first reaches perfect performance ($r = -0.63$, $p = 0.02$).

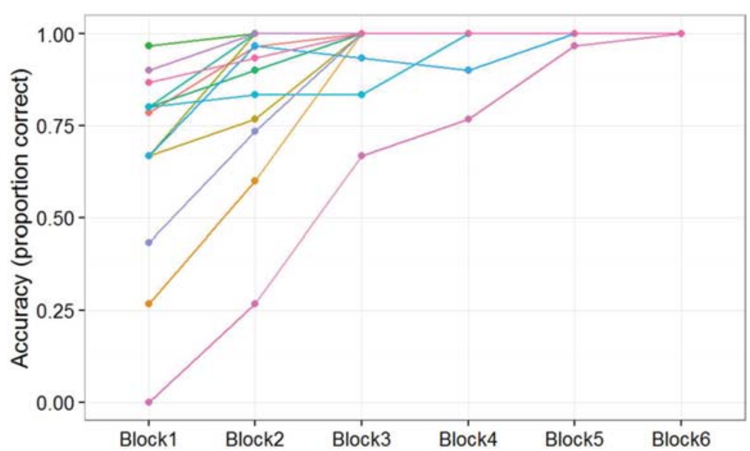


**Fig 4. Individual subject performance for knowledge* group.** Proportion of plays on which the subject scored as a function of block for subjects in the Knowledge* subset of the Predictable/No-Instruction group (subjects able to verbally report the predictive rule.) Each line shows a different subject from the group. All subjects in this group reached 100% performance in one of the blocks.

https://doi.org/10.1371/journal.pone.0179386.g004

## Trial of last error

To provide a more fine-grained look at what precedes the "perfect mastery", we examined subjects' performance on 5 trials just prior to each individual's trial of last error (TLE and TLE*; see [15], for an early study focusing on this measure in the context of concept learning). The TLE is specifically constrained to an individual's last trial such that an incorrect response was made. Hence, averaged across all subjects, the TLE has a mean accuracy of 0 with no variability. Given that the TLE is defined as each individual's last error in the experiment, all remaining trials must have perfect accuracy. The first TLE line (solid) shows performance hovering around 75%, followed (as must happen, given the definition of TLE) by a trial with zero accuracy (the last error) and then a performance of 100% on the remaining trials of the experiment. What is striking in the TLE is the relatively flat performance curve over this immediate pre-mastery period (see Fig 5), and the fact that the level of performance here (76%) is considerably higher than that seen overall in either the Control or the No-Knowledge* groups. (Superficially, the flatness of this figure might appear to paint a rather different picture of the buildup to insight than what is seen in Fig 3—but there is no conflict, because Fig 3 plots improvement over a far longer time-scale.) The TLE would suggest learners who ultimately attain mastery tend to reach that point by jumping up from a rather high plateau that already supports performance that is much better than what subjects in the Control and No-Knowledge* conditions generally ever attain.

However, after examining subjects' cursor locations during the course of Experiment 2, we came to think that our early termination criteria (a completed block of 100%) might make that TLE potentially a bit misleading. Suppose a subject gained insight to the predictive feature relatively early in a block. They could still make a motor error (e.g., not moving or clicking the mouse fast enough to score the point), and indeed they would have rather little incentive to

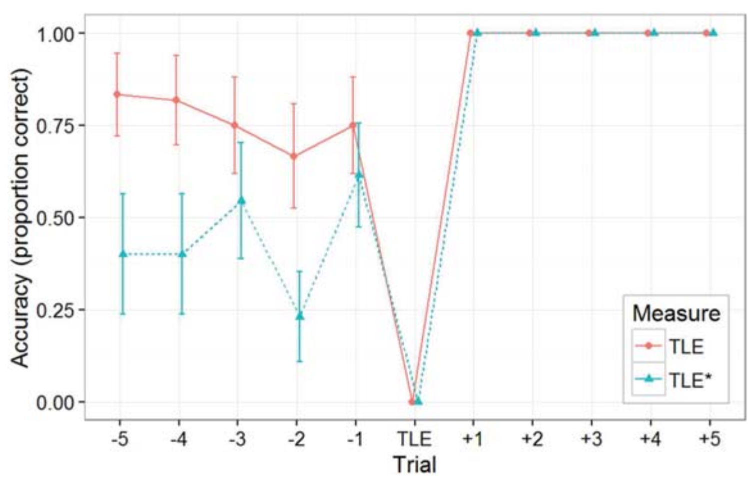


**Fig 5. Experiment 1 TLE and TLE* performance.** Performance for subjects in the Knowledge* subset of the Predictable/No-Instruction group for 24 trials prior to the subject's Trial of Last Error (TLE), averaging over pairs of adjacent plays. TLE(-5:-1) mean: 76%; TLE*(-5:-1) mean: 44%.

https://doi.org/10.1371/journal.pone.0179386.g005

score perfectly on the remaining trials within that block (since the reward of leaving early was available only for a perfect score on an entire block). We suspect this is likely the case given the non-perfect scoring in the Full Instructions group.

When subjects made a motor error after attaining insight, the apparent TLE would be shifted relative to the "true TLE". To deal with this subtle measurement pitfall, we computed another measure of the trial of last error, which we call TLE* (dashed). This is the last error that precedes five successive errorless trials (0.031 probability of occurrence by chance). The results of TLE* as rather convincing, in our opinion, in showing that the final mastery was accomplished *de novo* by jumping from a state of complete ignorance (44%) of where the demon will next appear.

## Experiment 2

To further investigate the surprising result where unawareness of the predictable feature led to performance that was no better than chance and final mastery was potentially de novo, we ran a higher powered version of the first experiment, this time specifically comparing Condition 1 (control) and Condition 2 (Predictive/No Instruction). In addition, we changed a few parameters to increase the difficulty of the game.

## Method

### Participants

Two hundred undergraduates from the same subject pool participated in this experiment for course credit. All were naïve to the purpose of the experiment.

### Stimuli

Stimuli were created using the same process as Experiment 1 with the following exceptions: Eye diameter determined direction of the demon {15px: Left, 30px: Right}. For the control condition, the eye diameter was randomly chosen from the discrete set {15px, 30px}. Eye color was randomly selected from the color set {yellow, blue} and assigned to 1 of 5 shades of the selected color. Horn height and width were drawn independently from a uniform distribution in the range [15 pixels, 60 pixels]. And lastly, the demon bodies were randomly assigned a color from the set {red, grey}.

### Procedure

The procedure was identical to Experiment 1 with the exception that Experiment 2 only included Condition 1 (Control) and Condition 2 (Predictive/No Instruction). In addition, the timing of a single trial was modified as follows: The V-shaped occluder appeared on screen alone for 2 seconds. Then the demon appeared at the bottom on the screen and paused for 3 seconds. After which, the demon would disappear behind the occluder for 2 seconds and reappear either or the left or right side of the occluder moving off screen with the demon being visible and clickable for 0.5 second. The trial concluded with a 2-second blank screen ISI.

## Results and discussion

A research assistant scored the hunches while being blind to the conditions. The first scoring criterion examined was that the subjects had to explicitly mention the predictive feature (i.e., the eyes). We also examined a stricter criterion where the statement had to include the idea of "big eyes to the right and small eyes to the left". No reliable differences between the two scoring criteria were observed, so we used the first scoring criterion for the data described here. Of the 100 subjects in the predictive/no-instruction condition, 40 subjects were counted as aware (Knowledge*) while 60 subjects were scored as No-Knowledge* (25 reported an incorrect hunch and 35 reported no hunch).

## Block performance

Fig 6 shows the same pattern of results. There was a significant main effect of group, $F(2, 197) = 186.7$, $p < 0.001$, and block, $F(5, 985) = 42.64$, $p < 0.001$, as well as a significant interaction, $F(10, 985) = 23.61$, $p < 0.001$. In line with Experiment 1, subjects who successfully described the predictive rule were able to reach a high level of performance while subjects that were unable to verbalize the predictive rule did not perform better than the control subjects, $F(1, 158) = 0.07$, $p = 0.79$. Subjects in the Knowledge* group did have more variability in performance (see Fig 7) in which some subjects were unable to complete at least one block at 100%. We would attribute this to the higher difficulty level of the game (i.e., faster demon exit speed).

## Trial of last error

The results from the TLE and TLE* analyses (see Fig 8) as described in Experiment 1 Results were strikingly similar to what was seen in Experiment 1. Performance prior to the TLE (solid) was relatively high (72%) given the stringent early termination criteria. However, as discussed above, the results of TLE* (dashed) indicate that the final mastery was again accomplished *de novo* by jumping from a state of complete ignorance (40%) of where the demon will next appear.

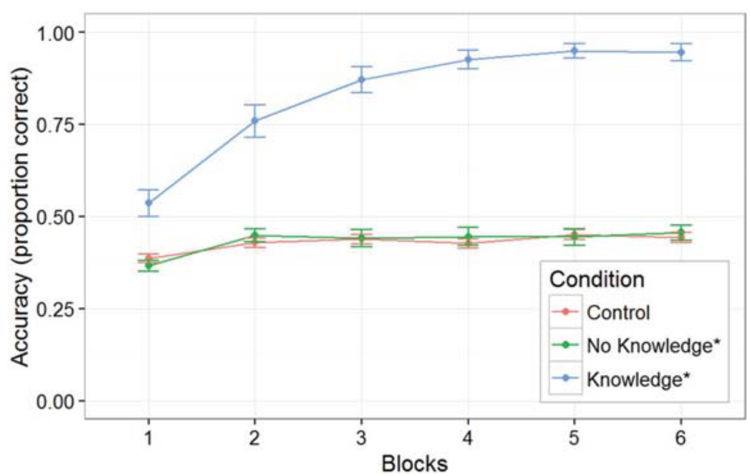


**Fig 6. Experiment 2 subject performance.** Proportion of plays resulting in successful scores for Control and Predictive/No-Instruction groups (Knowledge* subjects who were able to verbally report rule, and No-Knowledge* subjects who could not do so.) Averages were computed in the same manner as with Experiment 1. No-Knowledge* subjects scores were not superior to those of control subjects for whom there was no predictive relationship available to be exploited.

https://doi.org/10.1371/journal.pone.0179386.g006

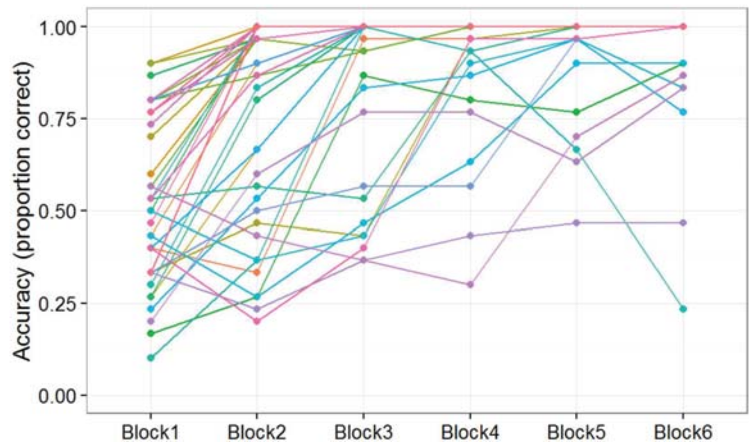


**Fig 7. Individual subject performance for knowledge* group.** Proportion of plays on which the subject scored as a function of block for subjects in the Knowledge* subset of the Predictable/No-Instruction group (subjects able to verbally report the predictive rule.) Each line shows a different subject from the group. The majority of subjects in this group reached 100% performance in one of the blocks.

https://doi.org/10.1371/journal.pone.0179386.g007

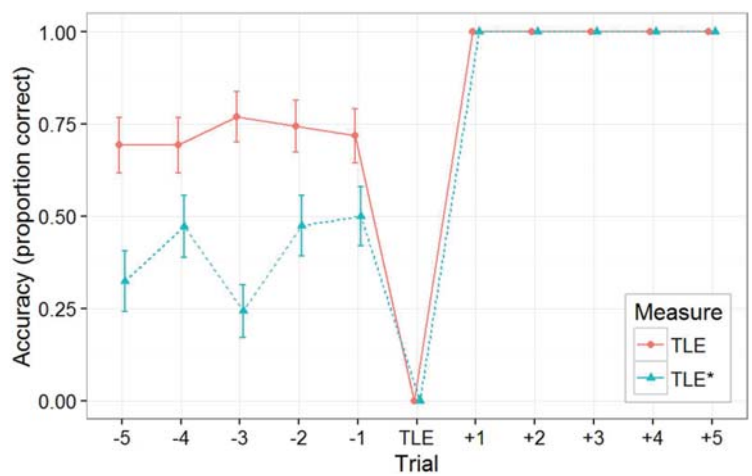


**Fig 8. Experiment 2 TLE and TLE* performance.** Performance for subjects in the Knowledge* subset of the Predictable/No-Instruction group for 12 trials prior to the subject's TLE, averaging over pairs of adjacent play. TLE(-5:-1) mean: 72%; TLE*(-5:-1) mean: 40%.

https://doi.org/10.1371/journal.pone.0179386.g008

## Cursor placement

To uncover the relationship between awareness and exploitation of the predictive feature in more fine grain, we also looked at where subjects placed the cursor just prior to the demon's exit of the occluder in each subject's last two blocks of plays. This gives us a more sensitive measure of subjects' predictions prior to making a committed response (i.e., the mouse click). Three possible ways subjects might play this game are to: (1) intentionally choose an exit based on a hunch, (2) randomly choose an exit, or (3) place the cursor in the middle to minimize the distance between the exits and the cursor. If subjects learned the predictive feature, either explicitly or implicitly, one would expect that they are more likely to use this to their advantage by placing the cursor at the correct tunnel exit before the demon exits the tunnel. Given the two distinctive strategies (side- and middle- choosing), the game screen was sectioned into equal thirds (Correct Side, Opposite Side, Middle) by pixels along the x-axis. For example, if a demon came out of the left tunnel, a mouse cursor on the left third of the screen would be labeled "Correct", a mouse cursor on the right third would be labeled "Opposite", and a mouse cursor in the middle third would be labeled "Middle". Fig 9 shows the proportion of cursor location by Group and cursor location for the last two blocks of game play. Across the 3 groups, cursor location was statistically significant, $F(2, 388) = 58.66$, $p < 0.001$.

For the Control group, there was a significant difference in where the mouse cursor was located, $F(2, 198) = 13.98$, $p < 0.001$. However, a post-hoc bonferroni-corrected pairwise t-test revealed subjects did not exhibit any correctly anticipate the side where the demon would

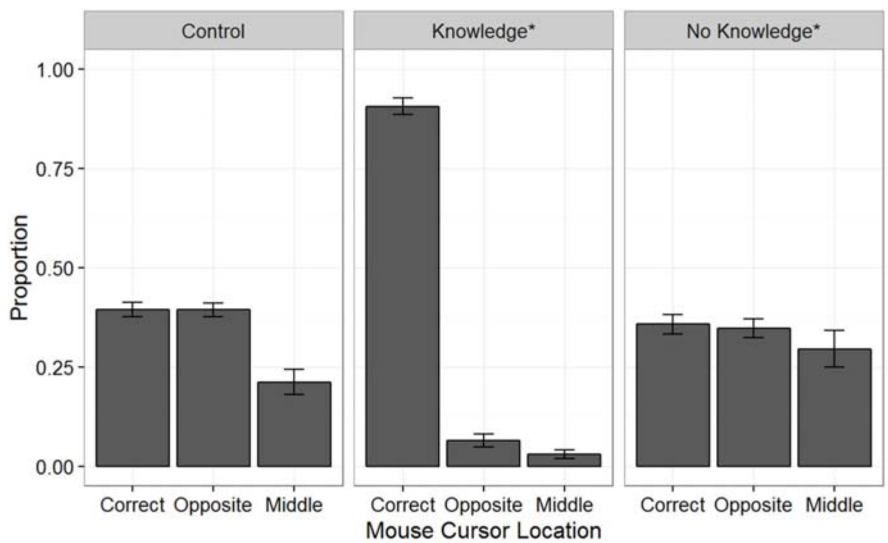


**Fig 9. Cursor locations.** Proportion of cursor locations just prior to the demon exit for the last two blocks of the game. The screen was divided into thirds: Correct (i.e., mouse cursor was on the side that the demon will exit), Opposite (i.e., the mouse cursor was on the opposite side that the demon will exit), and Middle (i.e., the mouse cursor was in the middle 1/3 region).

https://doi.org/10.1371/journal.pone.0179386.g009

appear (Correct vs. Opposite, $p = 1$). This was of no surprise and was to be expected given no predictive feature to be exploited.

For the Knowledge* group, there was a significant difference in where the mouse cursor was located, $F(2, 78) = 596.2$, $p < 0.001$. A post-hoc bonferroni-corrected pairwise t-test revealed significant differences for Correct vs. Opposite ($p < 0.001$) and Correct vs. Middle ($p < 0.001$) and no significant difference for Opposite vs. Middle ($p = 0.21$). In other words, Knowledge* subjects were able to exploit the predictive feature (i.e., small eyes and big eyes) by placing the cursor in the correct location (i.e., left and right, respectively) in anticipation of the exiting demon. In addition, we examined the cursor locations for the trials preceding TLE*. We would predict that if conscious insight did occur at the TLE*, the cursor locations for the preceding TLE* trials would not differ. This post-hoc analysis revealed no difference in cursor locations, $F(2, 76) = 1.86$, $p = 0.16$, suggesting there was no insight prior to the TLE*.

Lastly, subjects in the No-Knowledge* group did not seem to anticipate where the demons would exit, $F(2, 112) = 0.67$, $p = 0.51$. While these data (e.g., Correct vs. Middle, $p = 1$; Opposite vs. Middle, $p = 1$) show a different pattern compared to the Control group (e.g., Correct vs. Middle, $p = 0.001$; Opposite vs. Middle, $p = 0.001$), we suspect individual subjects used a combination of the side-guessing and neutral-region strategies to different degrees.

## General discussion

The studies reported here examined what happens when subjects play a simple video game that embodies a discrete and very useful predictive relationship (the height of a demon's horns, Exp. 1, or diameter of the demon's eyes, Exp. 2, predicts which way the demon will "choose" to go, information the player can exploit by laying in wait for the demon). Of most interest was what happened for a group of subjects who were not given any hints about the existence of the predictive regularity that they were exposed to (the Predictable/No-Instruction Condition). Based on the final exit-interview reports, this group of subjects could be divided surprisingly cleanly into two different subgroups: those who demonstrated clear (and verbalizable) conscious access to the rule (the Knowledge* subgroup), and those who did not (No-Knowledge* subgroup), with only a few subjects resisting easy classification.

Whereas past studies of implicit learning have generally found only a very weak or nonexistent relationship between conscious insight and behavioral indices of implicit learning, here the scores attained by the two groups very closely tracked their ability to verbally articulate the hidden rule. Indeed, for those showing no conscious access, there was no evidence that they acquired any ability whatsoever to make use of the predictive relationship in their game play. Their scores remained comparatively to the performance level of control subjects playing a random version of the game that did not allow prediction of where the demons would reappear. By contrast, the Knowledge* group, who were able to articulate the rule, showed dramatically better overall scores, almost all reaching the criterion of perfect performance within a whole block. In addition, Knowledge* subjects appear to accomplish this de novo and not from gradual learning (i.e., an abrupt jump in performance as well as a change in cursor placement behavior in pre-TLE* and post-TLE* trials).

While one might argue that accuracy is not sensitive enough show implicit learning and measurements such as reaction time must be used, our examination of the cursor data in Experiment 2 seem to shed some light on the No-Knowledge* subjects. Had there been any implicit learning, we would expect there to be at least a bias in where cursors were placed prior to any overt decision response of a mouse-click. Yet the pattern of nearly all No-Knowledge* subjects' cursor locations were in all three regions regardless of a demon's exit direction.

## Limitations and connections to previous findings

As noted above, the tight linkage between awareness and implicit usage of the regularity in the present study appears to conflict with the conclusion of several previous lines of research on implicit learning and awareness (e.g., [4,10,16]). Why is this? One possible reason is that the learning revealed in the earlier studies is encapsulated within the motor or perceptual systems. By contrast, in the game used in the current studies, useful learning based on the regularity would seem to require learning to base an action on the crucial predictor feature, a strategy change not just a modulation of response latency. Evidence for such unconscious learning is strikingly absent. One possibility, as noted earlier, is that the fact that the regularities embedded in Willingham et al. [4] and Miller [10] studies were only partially valid suppressed conscious utilization of these effects. Each of the differences noted suggests potentially testable hypotheses for follow-up research.

While the results challenge any assumption that implicit learning in general is unconscious, powerful, and "cognitively impenetrable", it is important not to overstate the conclusions. It is possible that with more training, non-conscious learning might have developed in this situation to take advantage of the predictive regularity embedded in our game. Perhaps this had even begun operating far too slowly to have produced any results that could be detected within the time limits of this study. Thus, it is possible that, like the proverbial tortoise, such a putative non-conscious learning process might eventually have caught up and enabled the player to exploit the predictive relationship without any conscious awareness of the relationship. This, too, is a testable interpretation.

A second limitation is that the results do not clearly indicate whether the learning process revealed here depends upon active and conscious hypothesis-testing, e.g., consciously and explicitly checking the relationship between each feature of the demon and the demon's behavior. It could be that the link to awareness exists because the learning arises from such a conscious reasoning process. Alternatively, as noted above, it could be that a slow buildup of information in a non-conscious learning system results in conscious awareness once a sufficiently high level of predictive success has been achieved. This kind of interpretation was suggested some years ago by Boakes in discussing the tight linkage between awareness and Pavlovian conditioning (see [17,18]). As mentioned above, the detailed time-course of learning here arguably fits somewhat better with the latter interpretation than it does with the conscious hypothesis-testing interpretation.

One intriguing aspect of the present results is the fact that subjects in the control condition often reported apparently illusory rules for the behavior of the demon (see Supplementary Online Materials) and in some cases, they voiced these with strong confidence. Future research might shed light on the interplay of conscious and non-conscious learning processes by examining in detail the stimulus displays seen by subjects who claimed to have discovered rules that are (at least on expectation) false. One question of interest is whether these reports reflect unusual local statistics of what these particular subjects actually experienced, or instead are completely fanciful.

A third limitation of the current study is that (quite by design) it used a predictive relationship that is potentially easy to verbalize. Naturally, predictive relationships that are not so easily verbalized (e.g., acquiring dexterity with motor skills like tennis playing and driving for which most people probably have only a poor descriptive vocabulary) may not show such a close connection to awareness (cf. [19]).

These limitations notwithstanding, the results of the present work lead to several conclusions. One is that it is possible to embed predictive regularities in computer games and to track their behavioral exploitation (comparing it to conscious reports), thus offering new ways to examine implicit learning using behavioral measures more compelling than modulation of reaction time. Second, given a highly reliable and useful predictive regularity in a game, the behavioral exploitation of this regularity can sometimes emerge with a far stronger linkage to conscious awareness than has generally been noted in the implicit learning field.

## Supporting information

**S1 File. Subjects' exit responses.**
(DOCX)

## Acknowledgments

This work was supported by the Office of Naval Research (MURI Grant #N00014-10-1-0072), a collaborative activity grant from the James S. McDonnell Foundation, and the National Science Foundation (Grant SBE-582 0542013 to the UCSD Temporal Dynamics of Learning Center and Grant SES 1461535 to PI: Michael Mozer). The opinions expressed here are those of the authors and do not represent views of the agencies that have supported this work. The authors thank Sorin Sodolescu for web programming assistance, to Noriko Coburn for excellent assistance with data analysis and presentation, and to Mike Mozer, Ray Perez, Tim Rickard for useful discussions. The authors also thank J. Bruce Morton for thoughtful comments on a previous version of this manuscript.

## Author Contributions

**Conceptualization:** HP.

**Data curation:** RT.

**Formal analysis:** RT.

**Funding acquisition:** HP.

**Investigation:** RT.

**Methodology:** HP RT.

**Project administration:** RT HP.

**Resources:** HP RT.

**Software:** RT HP.

**Supervision:** HP RT.

**Validation:** RT HP.

**Visualization:** RT.

**Writing – original draft:** HP RT.

**Writing – review & editing:** RT.

## References

1. Fu WT, Anderson JR. From recurrent choice to skill learning: A reinforcement-learning model. Journal of experimental psychology: General. 2006 May; 135(2):184.
2. Cleeremans A, Jiménez L. Implicit learning and consciousness: A graded, dynamic perspective. Implicit learning and consciousness. 2002:1–40.

3. Perruchet P, Vinter A. The self-organizing consciousness as an alternative model of the mind. Behavioral and Brain Sciences. 2002 Jun 1; 25(03):360–80.
4. Willingham DB, Nissen MJ, Bullemer P. On the development of procedural knowledge. Journal of experimental psychology: learning, memory, and cognition. 1989 Nov; 15(6):1047. PMID: 2530305
5. Curran T, Keele SW. Attentional and nonattentional forms of sequence learning. Journal of Experimental Psychology: Learning, Memory, and Cognition. 1993 Jan; 19(1):189.
6. Perruchet P. Learning from complex rule-governed environments: On the proper functions of nonconscious and conscious processes. Attention and performance XV: Conscious and nonconscious information processing. 1994;811–835.
7. Perruchet P, Gallego J. Association between conscious knowledge and performance in normal subjects: Reply to Cohen and Curran (1993) and Willingham, Greeley, and Bardone (1993). Journal of Experimental Psychology: Learning, Memory & Cognition. 1993; 19:1438–1444.
8. Willingham DB, Greeley T, Bardone AM. Dissociation in a serial response time task using a recognition measure: Comment on Perruchet and Amorim (1992). Journal of experimental psychology: learning, memory, and cognition. 1993 Nov 30; 19:1424–1430.
9. Reber AS. Implicit learning and tacit knowledge. Journal of experimental psychology: General. 1989 Sep; 118(3):219.
10. Miller J. Priming is not necessary for selective-attention failures: Semantic effects of unattended, unprimed letters. Attention, Perception, & Psychophysics. 1987 Sep 1; 41(5):419–34.
11. Miller J. The flanker compatibility effect as a function of visual angle, attentional focus, visual transients, and perceptual load: A search for boundary conditions. Attention, Perception, & Psychophysics. 1991 May 1; 49(3):270–88.
12. Hill T, Lewicki P, Czyzewska M, Schuller G. The role of learned inferential encoding rules in the perception of faces: Effects of nonconscious self-perpetuation of a bias. Journal of Experimental Social Psychology. 1990 Jul 31; 26(4):350–71.
13. Lewicki P, Hill T, Czyzewska M. Nonconscious acquisition of information. American psychologist. 1992 Jun; 47(6):796. PMID: 1616179
14. Hendrickx H, De Houwer J, Baeyens F, Eelen P, Van Avermaet E. Hidden covariation detection might be very hidden indeed. Journal of Experimental Psychology-Learning Memory and Cognition. 1997 Jan 1; 23(1):201–19.
15. Trabasso T, Bower G. Presolution reversal and dimensional shifts in concept identification. Journal of experimental psychology. 1964 Apr; 67(4):398.
16. Nissen MJ, Bullemer P. Attentional requirements of learning: Evidence from performance measures. Cognitive psychology. 1987 Jan 31; 19(1):1–32.
17. Boakes RA. How one might find evidence of conditioning in adult humans. Aversion, avoidance and anxiety: Perspectives on aversively motivated behavior. 1989:381–402.
18. Lovibond PF, Shanks DR. The role of awareness in Pavlovian conditioning: empirical evidence and theoretical implications. Journal of Experimental Psychology: Animal Behavior Processes. 2002 Jan; 28 (1):3. PMID: 11868231
19. Reed N, McLeod P, Dienes Z. Implicit knowledge and motor skill: What people who know how to catch don't know. Consciousness and cognition. 2010 Mar 31; 19(1):63–76. https://doi.org/10.1016/j.concog.2009.07.006 PMID: 19703779

Chapter 2, in full, is a reprint as it appears in Learning to exploit a hidden predictor in skill acquisition: Tight linkage to conscious awareness. Tran and Pashler (2017). The dissertation author was the primary investigatory and author of this paper.

# Chapter 3 **Relative proportion of exemplars and prototypes: What best promotes category learning?**

*Randy Tran and Harold Pashler*

**Abstract**

In a two-experiment study, we used a simple classification game to examine two aspects of perceptual category learning: 1) the classification of novel exemplars as a function of varying the relative number of exemplars and prototypes presented during training and 2) the linkage between awareness of category relevant features and individual subject performance. In the present study, participants learned to classify computer-generated “pumpkins” into two categories based on two features (size and hue). The ratio of exemplars to prototypes during training was varied (1:0, 0.75:0.25, 0.5:0.5, or 0.25:0.75) with a fixed 400 trial training. On the final test, participants were asked to categorize new exemplars generated from the same categories. Participants were then asked what (if any) strategies they used to classify the pumpkins. When the final test was given immediately after training, the groups that were trained with some prototypes performed higher than the group trained with only exemplars. There was, however, no group difference when the final test was given a week later. A strong linkage was found between individual subject’s reported strategies and their individual category boundaries. Subjects that were unaware of either one or both of the relevant features showed rule-based classification behavior instead of the correct information integration classification.

**Introduction**

Our visual systems have evolved to be very good pattern detectors, which in turn enable us to distinguish similarities and differences in our visual world. More importantly, we can distinguish group membership very quickly and without much cognitive effort. We are able to do this with not only objects we have seen before but also for novel objects. However, how are we able to identify something we have never seen before? The ability to identify and categorize objects in our visual world has had a long history of research with different views and theories, much of which have been thoroughly reviewed (e.g., Alfonso-Reese, 1996; Anderson, 1991; Homa, Rhoads, Chambliss, 1979; Posner & Keele, 1968; Maddox & Asbhy 1993; Nosofsky, Kruschke, & McKinley, 1992). For example, three prominent theories: Prototype theory (e.g., Homa et al. 1979; Posner & Keele, 1968; Smith & Minda, 1998), Exemplar theory (e.g., Kruschke, 1992; Nosofsky, 1988; Nosofsky, Kruschke, & McKinley, 1992), and Decision bound theory (e.g., Ashby & Maddox, 1993; Maddox & Ashby, 1993), have been used to describe our ability to classify visual objects. Many of these classic studies often have subjects learn categories that range from simple perceptual stimuli such as Gabor patches (e.g., Maddox, Ashby, & Bohil, 2003) and dot patterns (Posner & Keele, 1968) to more complex stimuli (e.g., Rabi, Miles, & Minda, 2015), usually to a degree of accuracy during training.

**Exemplars vs. Prototypes**

Given the many different category learning experimental manipulations (e.g., number of categories, linear vs non-linear boundaries, feedback vs. no-feedback, feedback timing, etc.) that have been used to examine the prototype and exemplary theories, we were surprised there has not been a study that manipulated prototype and exemplar exposure during training. This contrasts the categorization accuracy during training criteria that has been used in category learning. If the amount of exposure to category members plays a critical role in learning category structures, varying degrees of category exposure should affect how category boundaries are constructed. Thus, when creating a new mental category representation and decision boundaries, does having more exposure to exemplars or prototype instances promote better categorical classification of novel exemplars?

Although there are competing theories of how category members are identified, an individual's category representation is likely to consist of a mixture of exemplars and prototypes. Surprisingly, after carefully looking through the literature, we have found that previous paradigms do not directly manipulate the amount of exposure to exemplars and prototypes during training which a performance training criteria. We propose a more direct manipulation of category member exposure during training that will directly affect an individual's formulation of their mental category representations. It goes without saying that manipulating exemplar vs. prototype exposure should only produce a shift in the category representation. For example, individuals who are only exposed to exemplars during training are almost certainly abstracting categorical prototypes. However, it is likely that abstracting prototypes from exemplars introduces noise. Whereas those who have direct category prototype exposure will have a better categorical representation of category prototypes.

If people use prototypes to identify category members, then those who were exposed to only exemplars (with noisy abstracted prototypes) ought to perform worse on category identification compared to those who had direct prototype exposure.

**Individual Differences**

Recent findings in subject's explicit awareness and strategies used in classification have shown direct linkages between behavior (as measured by accuracy) and subjects' knowledge of category relevant attributes (e.g., Tran & Pashler, 2017) as well as defaulting to more parsimonious (i.e., simple over complex) strategies (e.g., Tran, Vul, & Pashler, 2017). Using our paradigm (described below), we also address the question: Are people consciously accessing exemplars/prototypes (or features of the category) during the classification of a new exemplar?

**Current Approach**

We had subjects play a game where they had to decide which family (A or B) different “pumpkin monsters” belongs. These pumpkins were computer generated stimuli in which the determinate categorical features were two dimensions that varied on a continuous scale. The conjunction of the two features (size and hue) were required for classification (i.e., diagonal decision boundary). The category prototypes were defined as the midpoint between the category boundary and the most extreme exemplars (i.e., the incenters of the two triangular category spaces).

During training, we devised four training schedules that manipulated the ratio of exemplars to prototypes (i.e., 1:0, 0.75:0.25, 0.50:0.50, 0.25:0.75) presented to subjects. By doing this between-subjects, we were able to directly control the exposure to exemplars and prototypes for each subject. For example, a subject in the 1:0 condition was exposed to only exemplars and no prototypes during training. On the other hand, a subject in the 0.50:0.50 condition would be exposed to exemplars for 50% of trials and see re-presentations of the two category prototypes for the other 50% of trials (all in a randomized order).

After training, subjects were given a final test with novel exemplars from the two categories either immediately (Exp. 1) or after a 1-week delay (Exp. 2). If subjects used exemplar representations of a category to classify novel objects, the conditions that have a higher exemplar to prototype ratio ought to have higher performances. On the other hand, if subjects used prototype representations of a category, we would expect the training condition with more prototype exposure to have a higher final test performance.

We also examined the classification of the two categories at the subject level. There has been greater interest in recent years to examine individual subject differences (Speelman & McGann, 2013) that may not be revealed by standard statistical practices (e.g., aggregating data across subjects). Given our previous work on awareness of relevant attributes (Tran & Pashler, 2017), we suspect that subject's explicit knowledge of the relevant category dimensions might play a role in classification boundaries and performance.

**Experiment 1**

Subjects were trained to classify stimuli into two categories in one of four conditions. The final test of novel exemplars was administered immediately after training.

**Method**

**Participants.** One hundred and thirty six undergraduates at the University of California, San Diego participated in this experiment for course credit. All were naïve to the purpose of the experiment. The subjects provided written informed consent, and the research was approved by the University of California San Diego Social and Behavioral Sciences Institutional Review Board.

**Stimuli.** The pumpkins differed on five features: square size (width/length in pixels), color of the body, left-eye radius, right-eye radius, and mouth width (see Figure 3.1 for examples). The width/height values were drawn from a uniform distribution of (200px, 400px). The color of the body was drawn from a uniform distribution hue with a range (0, 1/6) which is analogous to the range (0, 255) for green, 255 for red, and 0 for blue in RGB space. The radius of each eye was independently sampled from a Gaussian (mu = 25px, sigma = 7px). The mouth width was computed as W = (left-eye radius + right-eye radius) * (5 / 2) with a static height of 20px. The remaining facial features were static across all pumpkins: nose (each point was 20px from the center) and stem (W = 10px, H = 15px); these features were non-predictive of category assignment.

**Category Assignment.** The two categories (A and B) were determined by the size and hue of the pumpkins (i.e., a diagonal decision boundary). Due to size and hue being different units, proportions were used for category assignment. The prototype for category A was defined as the 75th percentile for size (i.e., 350px) and the 25th percentile for hue (i.e., 1/24) which is the point equidistant between the most extreme exemplar and the decision boundary. The prototype for category B was defined as the 25th percentile for size (i.e., 250px) and the 75th percentile for hue (i.e., 1/8). The three non-category relevant features varied for each prototype instance used in the experiment. Exemplars with size > hue were assigned to Category A and exemplars with size < hue were assigned to Category B. Figure 3.1 shows the category space for example stimuli for the prototypes and selected exemplars (exemplars on the decision boundary are for illustration purposes only).

**Design.** A four-level single-factor between-subjects design was used. Each condition consisted of two phases: training and testing. An exit survey was also given at the end of the testing phase. The training phase for each condition consisted of a different exemplar-to-prototype ratio. Condition 1 had a 1:0 ratio where subjects only saw exemplars for both categories during training and no prototypes. Condition 2 had a 0.75:25 ratio where 75% of the trials were of exemplars and 25% of trials were of prototypes. Condition 3 had a 0.50:0.50 ratio and Condition 4 had a 0.25:0.75 ratio. See Table 3.1 for the number of trials broken down by category type and ratio. The testing phase consisted of 96 novel exemplars (48 from Category A and 48 from Category B) as well as 96 trained items. The 96 trained items consisted of 24 exemplars from Category A and B (48 in total) and 24 prototype instances from Category A and B (48 in total).

**Procedure.** Subjects were tested individually in sound attenuated rooms for the computerized study. The entire session was completed in 1 hour.

**Training Phase.** All subjects in the randomly assigned four conditions read the same instructions:

> *In this experiment, you will be classifying pumpkins into 2 different families: the A-pumpkins and the B-pumpkins. You will have to figure out which family each pumpkin belongs to on your own. It may seem hard at first, but please try your best! You will get periodic rest periods during training. After training, there will be a final test on classifying the pumpkins into their respective families. Some of the pumpkins may look similar to each other, but please try your best to classify them into the correct families.*
>
> *Use the C-key to for the A-pumpkin family.*
>
> *Use the M-key for the B-pumpkin family.*
>
> *Please keep your index fingers on their respective keys for the duration of the experiment. Let your experimenter know if you have any questions at this time.*

Once subjects finished reading the instructions and had no questions, the experimenter began the training phase of the experiment by hitting the “Enter” key. Each trial showed a single pumpkin in the middle of the screen on a grey background with the question “Which family does this pumpkin monster belong to?” above the pumpkin. Subject responses were self-paced and feedback was given after a response was made. If subjects made a correct response, text feedback in green and a score counter displayed below the pumpkin (see Figure 3.2) and a correct tone would play. If subjects made an incorrect response, text feedback in red and a score counter displayed below the pumpkin and an incorrect tone would play. Feedback for correct responses displayed for 1-second and feedback for incorrect responses displayed for 2-seconds. The training phase consisted of four blocks with 100 trials each (a 20-second rest period was given after the first three blocks).

**Testing Phase.** Immediately following the last trial the fourth block, subjects were presented instructions for the testing phase of the experiment. Subjects completed testing phase without corrective feedback. The testing phase was 192 trials in length and consisted of 96 new exemplars generated from the two learned categories (i.e., 48 each) as well as 48 old exemplars (i.e., 24 each) and 48 old prototypes (i.e., 24 each) that were presented during training (condition 1:0 had 96 old exemplars given that no prototypes were seen during training). Old trained items were randomly intermixed with the new exemplars. Trials were self-paced and each trial began immediately after each response. A final score was presented on the screen after the last trial.

**Exit Survey.** After the completion of both training and testing phases, subjects were given an exit survey asking if they used any strategy in classifying the pumpkins into the two categories. Subsequently, all subjects were tested for color deficiency using the HRR pseudoisochromatic plates ($4^{th}$ ed.).

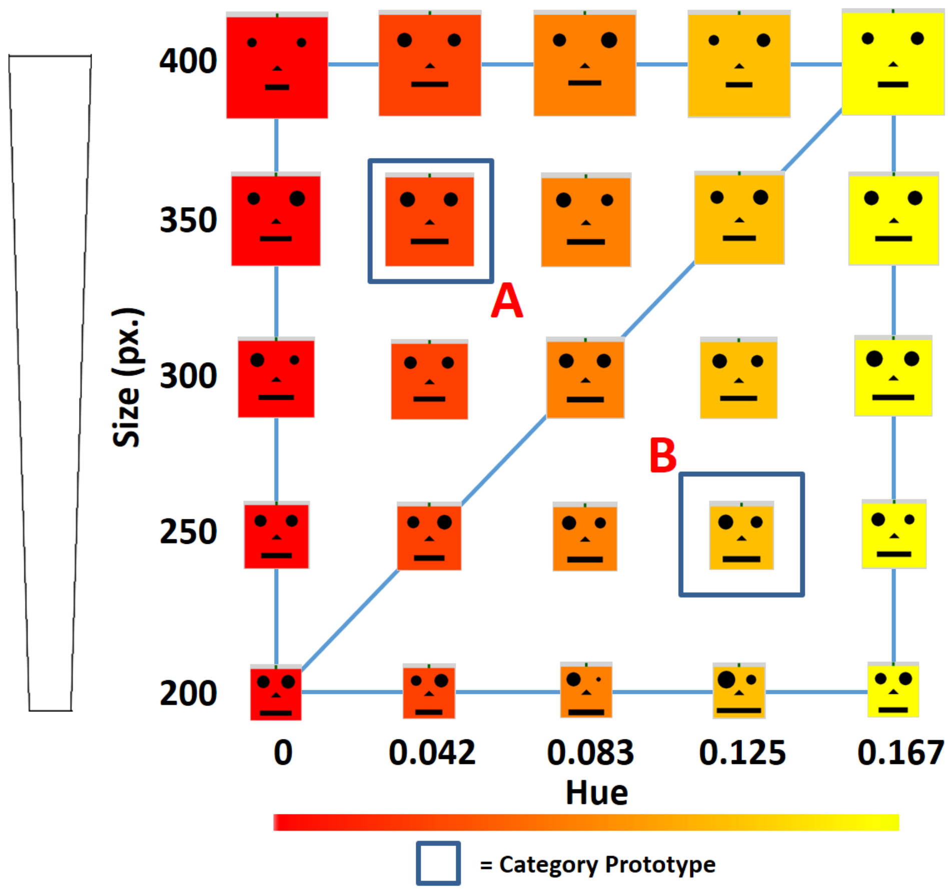


Figure 3.1. Stimuli category space. Individual pumpkins shown in this figure are examples of stimuli used in both Experiments 1 and 2 with both dimensions being continuous. Size ranges from 200 to 400 pixels for the width and height. Hue ranges from 0 to 1/6$^{th}$ to obtain a color spectrum between red and yellow. Values from the two dimensions were drawn uniformly from their respective ranges. Pumpkins were labeled Category A if the relative proportion for size was greater than the relative proportion for hue; for example, the prototype for Category A (top left boxed pumpkin) has a relative proportion 0.75 for size and a relative proportion of 0.25 for hue. Pumpkins were labeled Category B if the relative proportion for size was less than the relative proportion for hue. The lower right boxed pumpkin denotes the prototype for Category B (0.25 for size and 0.75 for hue). Pumpkins shown on the decision line are visual examples, these pumpkins were not used during training or testing.

Table 3.1. Trial breakdown for each condition in the training phase. There were a total of 400 trials during the training phase that consisted of different ratios of exemplars to prototypes. Subjects in Condition 1 (1:0) were only presented with exemplars during training; therefore, all 400 trials were exemplars split among the two categories (i.e., 200 trials of each category). Subjects in Condition 2 (0.75:0.25) were presented with 75% exemplars and 25% prototypes (i.e., 300 exemplar and 100 prototype trials). Subjects in Condition 3 (0.5:0.5) were presented with 50% exemplar and 50% prototypes (i.e., 200 exemplar and 200 prototype trials). Subjects in Condition 4 (0.25:0.75) were presented with 25% exemplar and 75% prototypes (i.e., 100 exemplar and 300 prototype trials). Note: the parameters for the category defining features on both prototypes (size and hue) were always the same. The 3 conditions with multiple prototype trials were representations of those prototypes (along with the other randomly varying features).

| **Condition** | **Training Exemplars** | **Training Prototypes** |
|---|---|---|
| **1 (1:0)** | A: 200<br>B: 200 | 0 |
| **2 (0.75:0.25)** | A: 150<br>B: 150 | A: 50<br>B: 50 |
| **3 (0.50:0.50)** | A: 100<br>B: 100 | A: 100<br>B: 100 |
| **4 (0.25:0.75)** | A: 50<br>B: 50 | A: 150<br>B: 150 |


Figure 3.2. A screenshot of the experiment during feedback after the subject successfully classified the pumpkin as Category A (training phase).

**Results and Discussion**

To be included in the analysis, subjects had to complete both phases of the experiment and the exit survey as well as pass the color blindness test. Ten subjects failed to complete the experiment and six subjects did not pass the color blindness test. The data from the remaining 120 subjects (30 subjects in each condition) were used for the analyses.

Training Performance. Figure 3.3 shows the learning performance for each block across the four conditions. There was a main effect of condition, $F(3, 464) = 116.77$, $p < 0.001$, and block, $F(3, 464) = 22.16$, $p < 0.001$, but no interaction, $F(9, 46) = 0.85$, $p = 0.57$.

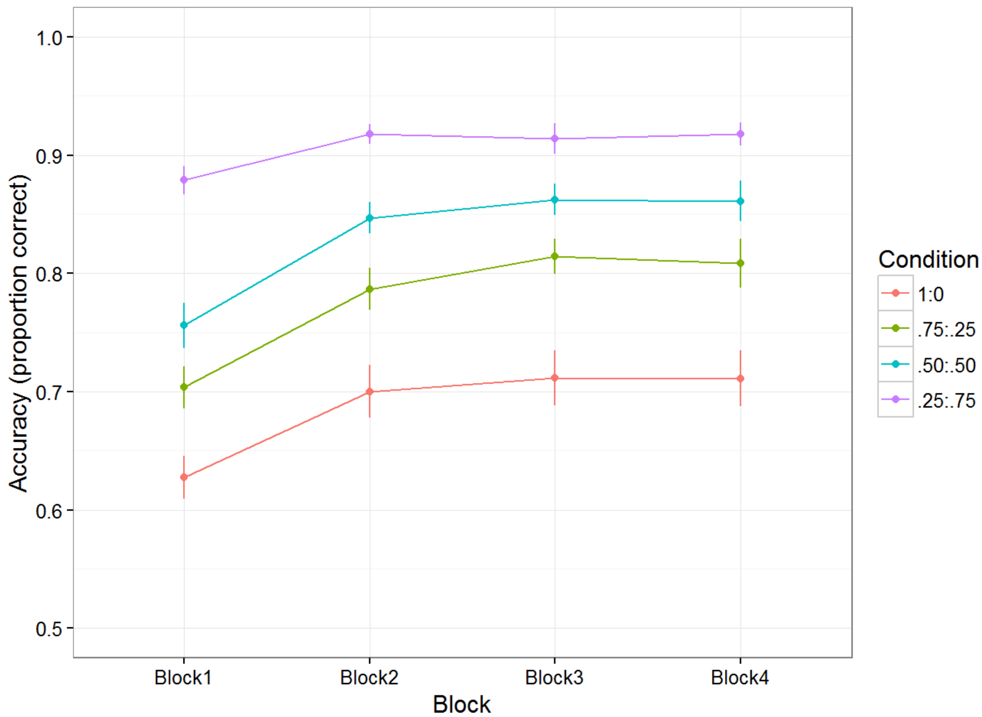


Figure 3.3. Training performance for Experiment 1. Each block consisted of 100 trials. Subject performance was averaged by condition.

Testing Performance. We examined the final test in two parts: recognition of previously presented exemplars/prototypes (Figure 3.4) and classification of novel exemplars. This separation allows us to critically examine how well subjects can classify new members of each category.

For the testing performance on novel exemplars, a one-way ANOVA revealed a significant difference among the means (1:0 $M$ = 0.719, $SD$ = 0.122, $SE$ = 0.022; 0.75:0.25 $M$ = 0.789, $SD$ = 0.091, $SE$ = 0.017; 0.50:0.50 $M$ = 0.764, $SD$ = 0.085, $SE$ = 0.015; 0.25:0.75 $M$ = 0.769, $SD$ = 0.075, $SE$ = 0.014) across the four conditions, $F(3, 116) = 2.921$, $p = 0.037$ (see Figure 3.5). All conditions had performance statistically above 50% chance ($ps < 0.001$). The three conditions that included prototypes during training did not differ on the final test: 0.75:0.25 vs. 0.50:0.50, $t(58) = 1.12$, $p = 0.269$; 0.75:0:25 vs.

0:25:0.75, $t(58) = 0.918$, $p = 0.362$; 0.50:0.50 vs 0.25:0.75, $t(58) = 0.268$, $p = 0.789$. This suggests that while any combination of exemplars to prototype ratio promotes learning of the categories on an immediate test with new category exemplars, subjects exposed to at least some prototypes during training had better category representations (as inferred by their final test performance).

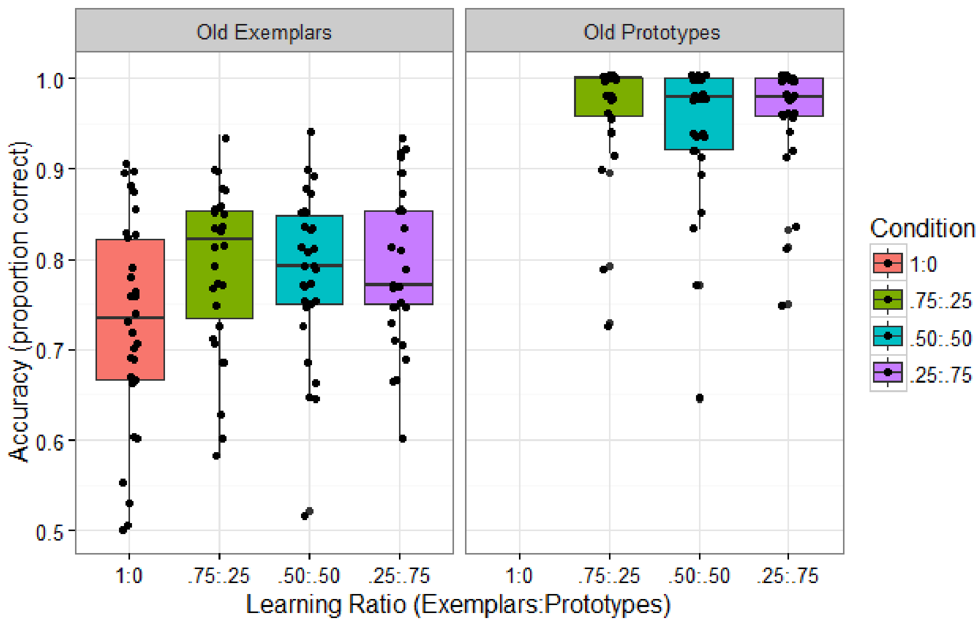


Figure 3.4. Experiment 1 boxplots for final test performance on previously seen exemplars and prototypes. Note: No prototypes were seen in Condition 1.

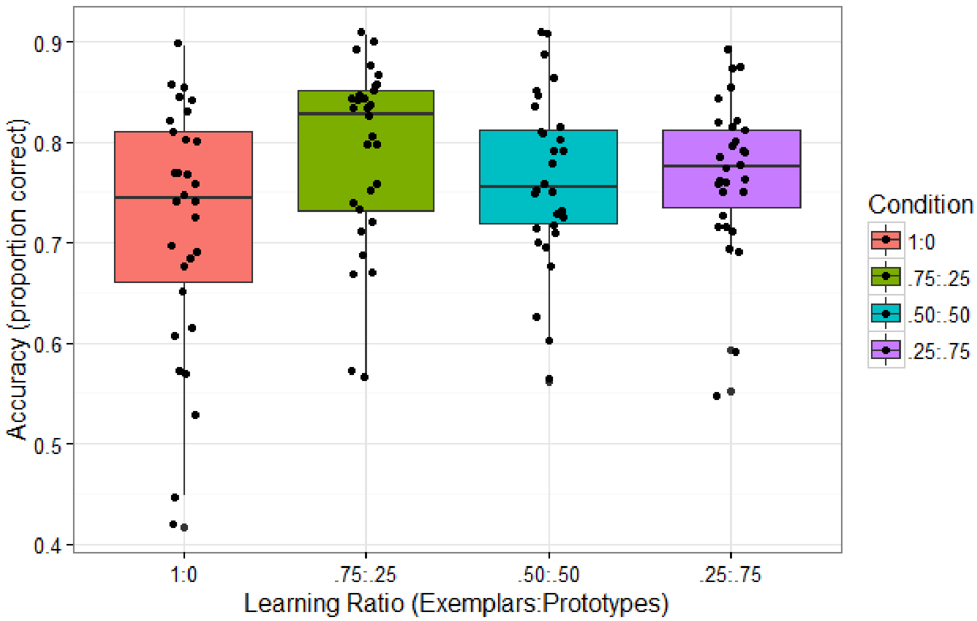


Figure 3.5. Experiment 1 boxplots for final test performance on New Exemplars by Condition.

In addition to group means, we also examined individual subject performance. Instead of looking at group means or even individual means, we visualized subjects' data by plotting their category space heat maps for classification responses on the final test for novel exemplars. Figure 3.6 shows an example of 24 subjects' (6 from each condition) individual category space heat maps. Each of the tiny color squares represent proportion of trials classified as Category A (i.e., more A responses in the binned area are more red and less A responses in the binned area are more purple; less A responses effectively means more B responses were made). Each individual's heat map maps onto the category space depicted in Figure 3.1 where the X-axis is the hue and the Y-axis is the size. The ideal observer with perfect classification ability would have a diagonal decision boundary (Figure 3.7, left) where the upper left region would be all red and lower right region would be all purple. Clearly not all subjects show this pattern of results. In fact, there are numerous subjects with vertical- and horizontal-like boundaries, which suggests these subjects are using only one of the two critical parameters. The vertical-like boundaries (more red on left and more purple on right) are likely due to subjects using only hue (Figure 3.7, middle) and horizontal-like boundaries (more red on top and more purple on bottom) are likely due to subjects using only size to classify the exemplars on the final test.

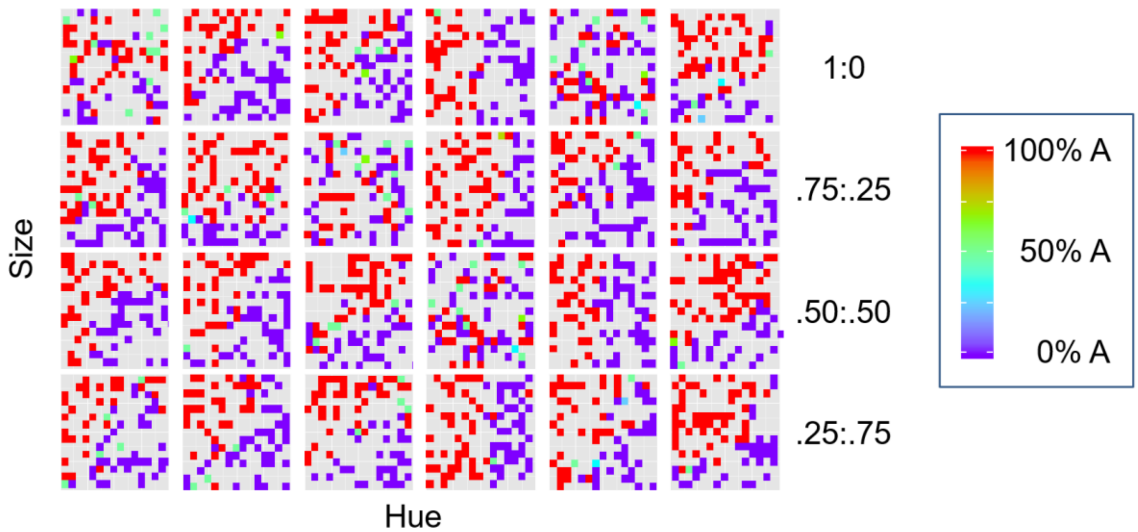


Figure 3.6. Item level data from 24 selected subjects. Each box depicts an individual subject's classification behavior. The color of the tiny squares within each box represents the proportion of trials that the subject classified as Category A, as a function of hue (horizontal axis) and size (vertical axis).

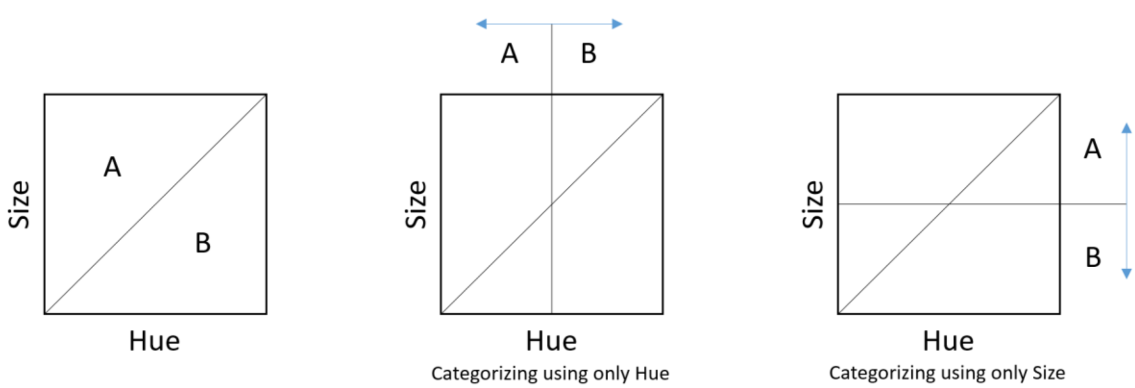


Figure 3.7. *Left*: perfect categorization. *Middle*: categorization based on hue. *Right*: Categorization based on size.

While our visual analysis clearly showed our subjects engaging in different classification strategies, a more in depth analysis is required to determine whether subjects actually learned the true classification boundary and if their classification was through an explicit strategy.

A research assistant, blind to the conditions, scored each subject's exit survey for their reported classification strategy. For example, if a subject stated on the exit survey that they used color (and irrelevant features such as eye size or mouth width) to determine the category of the pumpkins, that subject was labeled as a Color Only strategy user. Subjects that mention size (and irrelevant features) were labeled as Size Only strategy users. Those that state color and size (and irrelevant features) in their exit survey were labeled as Color AND Size users. Finally, subjects that did not mention either color or size were labeled as Other (this consisted of subjects that provided no text on the exit survey or stated irrelevant features). A logistic regression was used on each subject's final test classification. The slope of the best fit from the model was extracted for each subject and converted into an angle using atan(slope) * 180/pi for a possible range of angles [0, 180]. For example, the ideal observer with a perfect decision boundary would have a fitted regression with a slope equal to 1. The corresponding angle would 45 degrees. Therefore, subjects that used the correct integration of both color and size should have angles around 45. Subjects that used only color would have a vertical best fit line which translate to angles near 90 degrees and subjects that used only size would have a horizontal best fit line which translates to angles near 0 degrees. Lastly, using neither relevant features should result in a uniform distribution in angle space. The distributions for the four reported strategies are plotted in Figure 3.8. There are clearly four distinct distributions such that subjects' reported strategies map on directly to their classification behaviors. To test the equality of the probability distributions, the Two-sample Kolmogorov-Smirnov test was used. All three distributions were statistically different from each other, Color & Size vs Size Only: $D = 0.73$, $p < 0.001$, Color & Size vs Color Only: $D = 0.68$, $p < 0.001$, Color Only vs Size Only: $D = 0.84$, $p < 0.001$. The subjects that used Color Only, Size Only, and Other features appear to show no learning of the information-integration decision boundary.

While the ANOVA on group level means (as shown above) would lead us to believe that there is clear learning of the category boundary and ability to classify new exemplars in every condition (i.e., all above chance), examining individual subject data depicts a different story. Only a subset of subjects (across all conditions) learned the relevant features to classify the stimuli (the integration of color and size). More importantly, subjects' responses reflected that of their behavioral responses suggesting clear explicit strategy in classification. In addition, many subjects used the simple rule-based strategies (i.e., Color Only or Size Only) rather than the information-integration strategy (i.e., Color and Size) consistently through the experiment (see Figure 3.9).

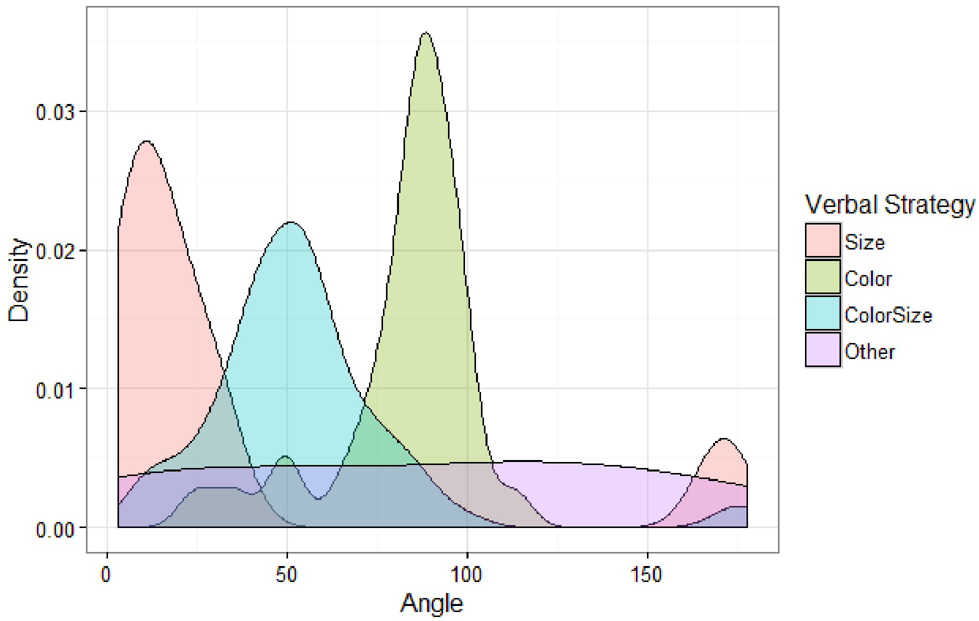


Figure 3.8. Distribution of classification angles by the four possible exit strategy (as determined by an RA blind to the conditions) for Experiment 1. Possible angles range from 0 to 180. Subjects that stated they used color and size to classify the novel exemplars have a distribution centered near the ideal 45 degrees. Subjects that stated they used color but not size have a distribution centered near the ideal 90 degrees. Subjects that stated they used size only but not color have a distribution centered near the ideal 0 degrees (because the angles are bounded between 0 and 180, a transformation of a slight negative slope from the regression line results in a large angle value—hence the distribution bump for angle values 150 and greater for size). Lastly, subjects that did not provide a strategy or stated irrelevant features had a uniform distribution of angles.

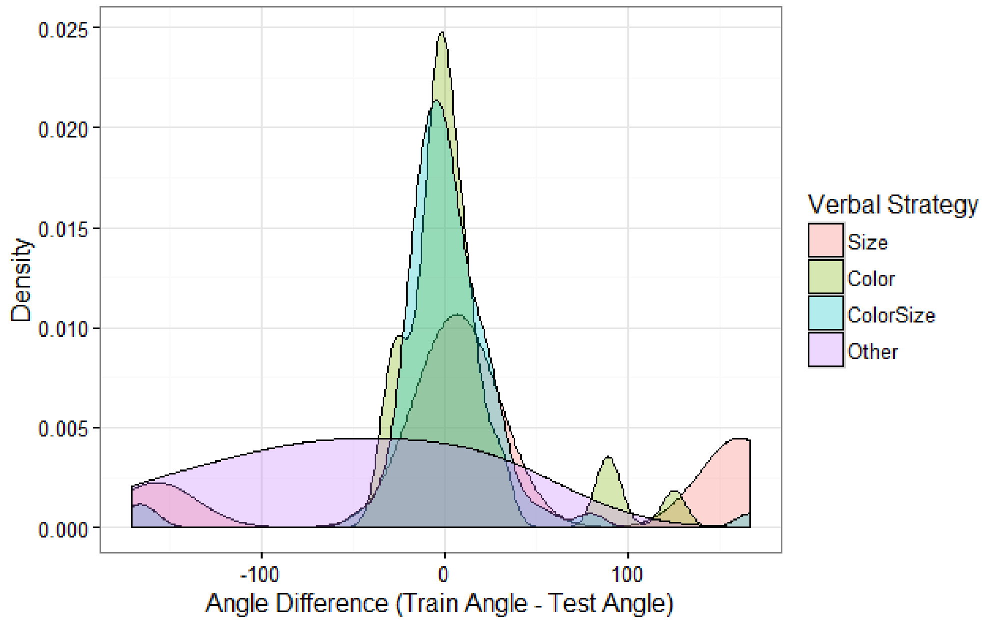


Figure 3.9. Distributions for the classification angle differences between final block during training and final test (Experiment 1). Possible angles range from -180 to +180. All four reported strategies essentially center around 0. In other words, subjects' explicit reported strategies on the final test were also utilized on the last block of training. For example, a subject that used both size and color to classify the pumpkins during both training and testing would have angles of 45 degrees, thereby resulting in a difference of 0.

## Experiment 2

Experiment 1 showed that having some exposure to prototypes during training promoted better learning of the category space. In our second experiment, we examined whether this training advantage would still exist if the final test was given at a delay. Here we examine the fidelity of learning by using the same experimental paradigm with an added delay of 1-week between training and testing.

## Method

**Participants.** Two hundred and seventy one undergraduates from the same subject pool participated in this experiment for course credit. All were naïve to the purpose of the experiment.

**Stimuli.** Stimuli were created using the same process as Experiment 1.

**Design.** The design was identical to Experiment 1 with the following exception: the testing phase consisted of only novel exemplars (100 from Category A and 100 from Category B). The old exemplars and prototypes were removed from the final test to avoid any possible influence on the classification for novel exemplars.

**Procedures.** The procedures were identical to Experiment 1 with the exception that Experiment 2 had a 1-week delay between training and testing.

**Results and Discussion**

As with Experiment 1, to be included in the analysis, subjects had to complete both phases of the experiment and the exit survey as well as pass the color blindness test. One subject was excluded for being off task, five subjects were excluded for failing the color blindness test, three subjects were excluded for not completing the exit survey, and 22 subjects were excluded due to improper data logging by a research assistant. The data from the remaining 240 subjects were used in these analyses.

Training Performance. The performance for training across the four conditions mirror Experiment 1 (see Figure 3.10). There was a main effect of condition, $F(3, 944) = 300.249$, $p < 0.001$, and block, $F(3, 944) = 49.445$, $p < 0.001$, but no interaction, $F(9, 944) = 1.521$, $p = 0.136$.

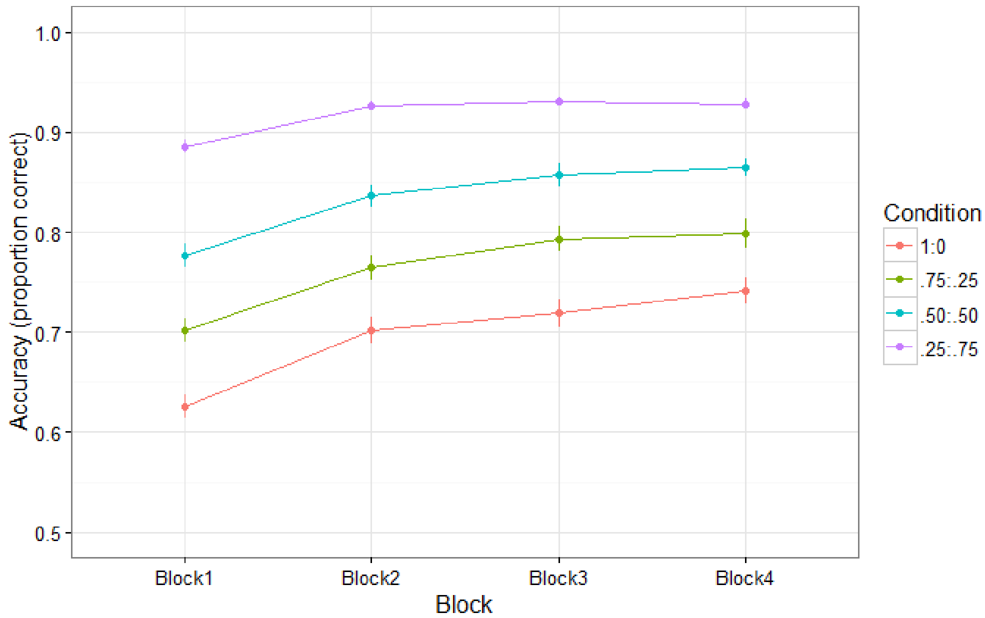


Figure 3.10. Training performance for Experiment 2. Each block consisted of 100 trials. Subject performance was averaged by condition.

Testing Performance. A one-way ANOVA revealed no significant difference among the means (1:0 $M$ = 0.693, $SD$ = 0.119, $SE$ = 0.015; 0.75:0.25 $M$ = 0.699, $SD$ = 0.153, $SE$ = 0.020; 0.50:0.50 $M$ = 0.689, $SD$ = 0.155, $SE$ = 0.020; 0.25:0.75 $M$ = 0.691, $SD$ = 0.162, $SE$ = 0.021) across the four conditions, $F(3, 236) = 0.047$, $p = 0.987$ (see Figure 3.11). Subjects in all conditions performed statistically above 50% chance ($ps < 0.001$). In contrast to Experiment 1, subjects who were given the final test after a 1-week delay showed no benefit of having some prototype exposure during training.

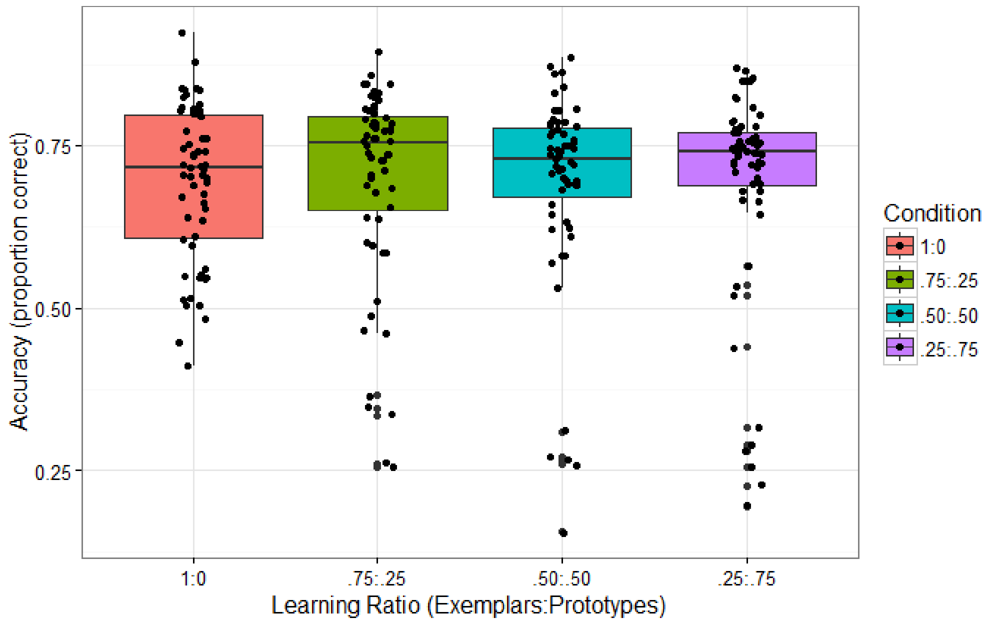


Figure 3.11. Experiment 2 boxplots for final test performance on New Exemplars by Condition.

Individual Subject Performance. Using the same method as Experiment 1, we first visualized subjects' performance with a category space heat map. 24 subjects from the four conditions are shown in Figure 3.12. Similarly to Experiment 1, there were subjects showed a variety of category classification. We used the same analysis procedure as described in Experiment 1 (i.e., grouped subjects by stated exit strategy and extracted classification boundary angle from logistic regression). The distributions in Figure 3.13 and Figure 3.14 show an identical pattern to that of Experiment 1. Although the ANOVA showed that the group means were not statistically different from each other, we still found, within each condition, numerous subjects using different categorization strategies. The three main distributions were statistically different from each other, Color & Size vs Size Only: $D$ = 0.66, $p < 0.001$, Color & Size vs Color Only: $D = 0.52$, $p < 0.001$, Color Only vs Size Only:

*D* = 0.78, *p* < 0.001. Subjects that did not clearly state an explicit strategy or mention the relevant attributes showed a similar uniform-like distribution to that of Experiment 1.

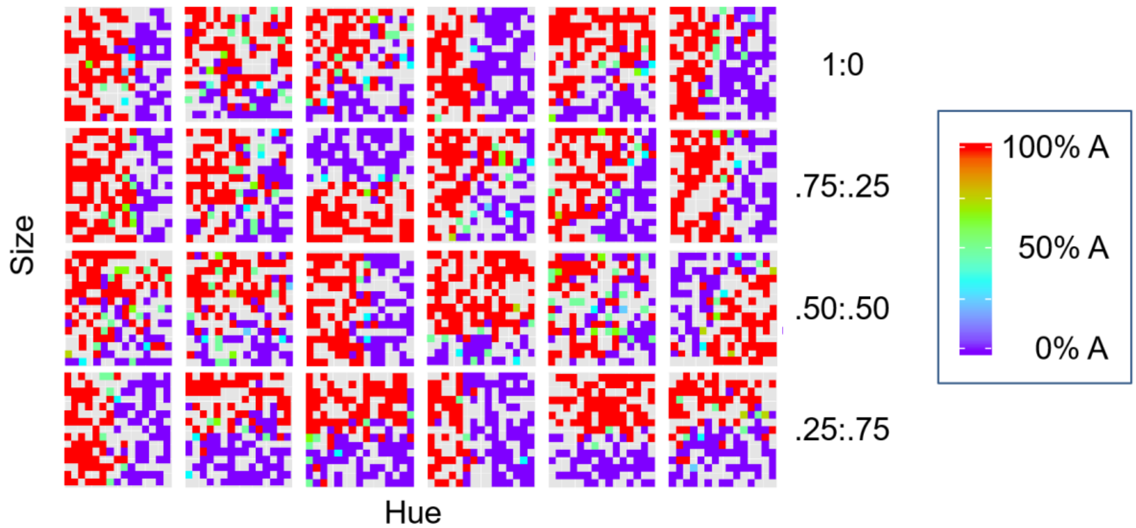


Figure 3.12. 24 selected subject from their respective training conditions for Experiment 2. Each box depicts a subject's respective final test response for novel exemplars contingent on Category A. A perfect final test would have a positive-slope diagonal boundary such that all response above are colored red and all responses below are colored purple. Cells are denser due to more final test trials in Experiment 2.

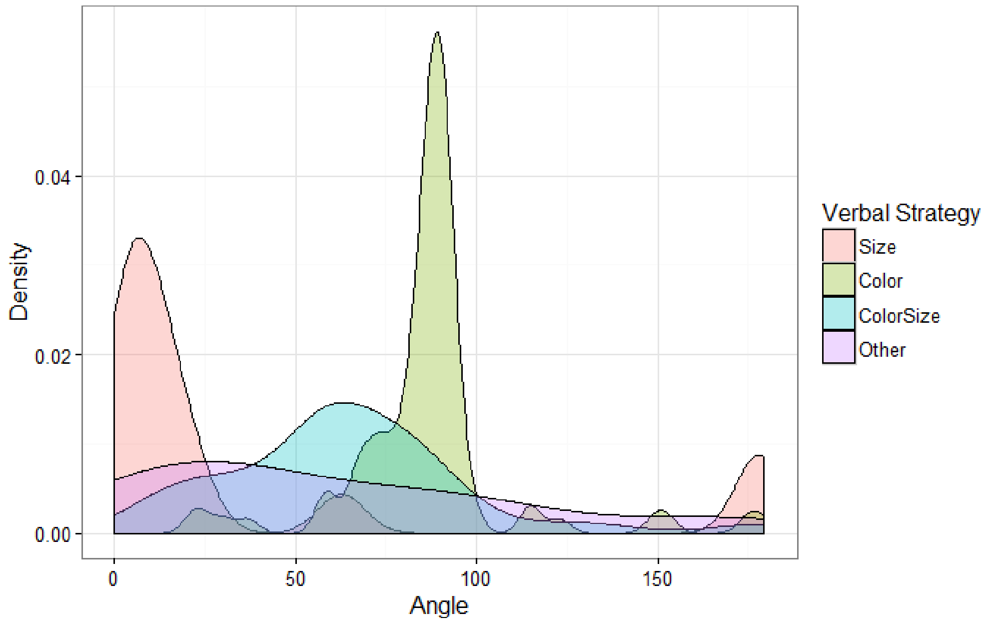


Figure 3.13. Distribution of classification angles by the four possible exit strategy (as determined by an RA blind to the conditions) for Experiment 2. Possible angles range from 0 to 180. Subjects that stated they used color and size to classify the novel exemplars have a distribution centered near the ideal 45 degrees. Subjects that stated they used color but not size have a distribution centered near the ideal 90 degrees. Subjects that stated they used size only but not color have a distribution centered near the ideal 0 degrees (because the angles are bounded between 0 and 180, a transformation of a slight negative slope from the regression line results in a large angle value—hence the distribution bump for angle values 150 and greater for size). Lastly, subjects that did not provide a strategy or stated irrelevant features had a uniform distribution of angles.

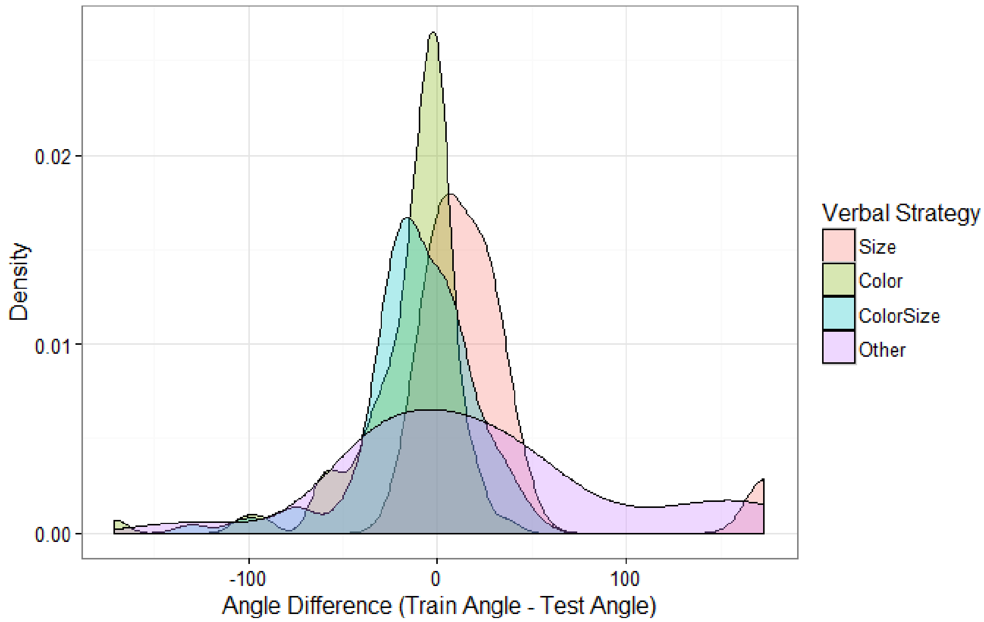


Figure 3.14. Distributions for the classification angle differences between final block during training and final test (Experiment 2). Possible angles range from -180 to +180. All four reported strategies essentially center around 0. In other words, subjects' explicit reported strategies on the final test were also utilized on the last block of training. For example, a subject that used both size and color to classify the pumpkins during both training and testing would have angles of 45 degrees, thereby resulting in a difference of 0.

## General Discussion

In Experiment 1, we found that subjects benefited from some exposure to at least some prototypes during training. Subjects were better able to classify novel exemplars into correct categories. However, this effect disappears after a one-week delay as seen in the final test performance for Experiment 2. Although we set out to experimentally manipulate how subjects construct their mental representations for each category, our in depth individual-subjects analyses revealed a different pattern of results. The aggregate means are misleading. Not all subjects showed learning of the diagonal decision boundary by integrating both relevant dimensions (size and color). In fact, there were many subjects that used a unidimensional boundary as seen by vertical and horizontal decision boundaries.

The studies reported here examined if subjects' categorization ability are affected by different levels of exemplar and prototype exposure during training. Our first experiment revealed a small effect where some prototype exposure during training promoted higher performance on classification of new exemplars. However, this effect disappeared when the final test was given at a one-week delay as seen in Experiment 2.

More importantly, while the aggregate data showed relatively good classification ability by subjects, our individual subject analysis revealed that many subjects did not learn the correct classification boundary. To correctly classify a given exemplar, subjects must integrate both the color and size of the stimuli; instead, subjects used only one of the two attributes. Taken altogether, our studies show that there may be a small effect of having more prototype exposure during training; however, this disappears when classification of new exemplars is not immediate. Why might this be? Given our individual subject analysis, we speculate that some subjects may have learned the relevant features (color and size) but forgot the integration after a week and only used one of the two relevant features.

Chapter 3, is currently being prepared for submission for publication of the material. Tran and Pashler. "Relative proportion of exemplars and prototypes: What best promotes category learning?" The dissertation author was the principal researcher and author of this material.

**References**

Alfonso-Reese, L. (1996). Dynamics of category learning. Unpublished PhD Dissertation. University of California, Santa Barbara.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409–429. https://doi.org/http://dx.doi.org/10.1037/0033-295X.98.3.409

Ashby, F. G., & Maddox, W. T. (1993). Relations between Prototype, Exemplar, and Decision Bound Models of Categorization. *Journal of Mathematical Psychology*, *37*(3), 372–400. https://doi.org/10.1006/jmps.1993.1023

Homa, D., Rhoads, D., & Chambliss, D. (1979). Evolution of conceptual structure. *Journal of Experimental Psychology: Human Learning and Memory*, *5*(1), 11–23. https://doi.org/http://dx.doi.org/10.1037/0278-7393.5.1.11

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*(1), 22–44. https://doi.org/http://dx.doi.org/10.1037/0033-295X.99.1.22

Maddox, W. T., & Ashby, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, *53*(1), 49–70. https://doi.org/10.3758/BF03211715

Maddox, W. T., Ashby, F. G., & Bohil, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(4), 650–662. https://doi.org/http://dx.doi.org/10.1037/0278-7393.29.4.650

Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(4), 700–708. https://doi.org/http://dx.doi.org/10.1037/0278-7393.14.4.700

Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(2), 211–233. https://doi.org/http://dx.doi.org/10.1037/0278-7393.18.2.211

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*(3, Pt.1), 353–363. https://doi.org/http://dx.doi.org/10.1037/h0025953

Rabi, R., Miles, S. J., & Minda, J. P. (2015). Learning categories via rules and similarity: Comparing adults and children. *Journal of Experimental Child Psychology*, *131*(Supplement C), 149–169. https://doi.org/10.1016/j.jecp.2014.10.007

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(6), 1411–1436. https://doi.org/http://dx.doi.org/10.1037/0278-7393.24.6.1411

Speelman, C., & McGann, M. (2013). How Mean is the Mean? *Frontiers in Psychology*, *4*. https://doi.org/10.3389/fpsyg.2013.00451

Tran, R, & Pashler, H (2017). Learning to exploit a hidden predictor in skill acquisition: Tight linkage to conscious awareness. *PLOS ONE, 12*(6): e0179386. https://doi.org/10.1371/journal.pone.0179386

Tran, R., Vul, E., & Pashler, H. (2017). How effective is incidental learning of the shape of probability distributions? *Royal Society Open Science*, *4*(8), 170270. https://doi.org/10.1098/rsos.170270

## Conclusion

My dissertation addressed the following three questions: 1) Does exploitation of predictive relationships occur when there is no explicit knowledge of the underlying structure? 2) Might some of the findings in implicit learning have very limited generalizability? 3) Lastly, does averaging across subject data mask what is learned by individuals? Throughout my dissertation, I demonstrated how explicit awareness can affect the learning of predictive relationships between task-relevant features and how simpler more discrete components were utilized within participants' strategies. In Chapter 1, we explored whether participants could induce a bimodal probability distribution through a whack-a-mole game. It was only when the bimodality was discretized did participants show any learned behavior of a bimodal distribution. Extending these findings, we demonstrated in Chapter 2, participants' used the simple rule to their advantage (e.g., placing their cursor at the correct tunnel exit) only when they became explicitly aware of the rule. Participants that were unable to verbalize the task-relevant feature were unable to utilize the task relevant feature and instead relied on a variety of ill-formed strategies. In Chapter 3, we explored a more complex predictive relationship. Similarly to Chapter 2, participants relied heavily on their explicit knowledge to formulate their classification strategies.

In the experiments described throughout all three chapters, participants' behaviors favored simpler components of complex rule that closely tracked their explicit knowledge. While prior research has shown implicit learning occurring (e.g., with amnesiac patients), it is very likely that not all learning takes place implicitly and even tasks once believed to be done implicitly may be done in an explicit manner. Specifically, we found in Chapters

2 and 3 that deliberate and conscious formation of strategies played a big role in learning where nearly all participants showed awareness of the task-relevant features.

While these findings taken altogether suggest exploitation of predictive relationships are done explicitly, follow-up studies will be needed to determine boundary conditions in which these occur. For example, an exit survey follow-up for Chapter 1 is needed to directly compare the impact of continuous and discrete (Chapters 2 and 3) responses. This would give us more insight into whether discreteness of the to-be-measured response influences participants' strategies (e.g., Chapter 1 participants may have been aware of the bimodal distribution but the continuous nature of the task masked their knowledge). The findings from Chapters 2 and 3 suggest complexity of the to-be-learned features play a role in participants' awareness and strategies. A logical follow-up to Chapters 2 and 3 would be to use features that use probabilistic outcomes. This would address whether all-or-none processes are used and how much insight participants are able to gain.

In summary, my dissertation work examined how people exploit predictive relationships to their advantage where there is a strong tendency to favor simple and verbalizable relationships and are used explicitly for optimal performance.