ENCODING TEMPORAL ORDER AND VISUAL STATISTICAL LEARNING

By

WONN SANG PYON

A Thesis Submitted to The Honors College

In Partial Fulfillment of the Bachelor's degree
With Honors in
Neuroscience and Cognitive Science

THE UNIVERSITY OF ARIZONA

M A Y 2 0 1 5

Approved by:

Dr. Rebecca Gómez
Psychology

Abstract

The literature suggests that visual statistical learning occurs from a very early age, with evidence suggesting that newborns are able to discern between familiar and novel sequences at just 2 days old. However, based on recent findings on the role of the medial temporal lobe in visual statistical learning in combination with our current understanding of this region's developmental timeline, we believe children younger than 40-months are unable to discern between the temporal regularities found between shapes in a sequence. In this particular study, we piloted two learning paradigms on adult subjects expecting to see a clear ability for the adult subjects to discriminate between our three categories of temporal order. Performance for our first paradigm, Fade-to-Reveal, revealed a significant improvement in reaction times through training, indicative of learning. For our second learning task Search-and-Find, the results of training suggested initial improvement with a regression in performance due to fatigue. Interestingly, subjects for both paradigms showed no real ability to explicitly recall the different shape-pairs at test. We interpret these opposing results to indicate that learning in these paradigms is implicit and thus the explicit recall test is not an appropriate measure of knowledge on shape-pairs.

Introduction

Statistical learning is a generalized learning mechanism we all utilize. It allows us to perceive patterns and regularities in the world around us and it is understood to be a critical part of our development from a very early age. Eight-month-old infants are able to utilize this particular type of learning effectively to rapidly learn the rules hidden in our languages and it is reported that infants were able to process and extract statistical regularities given just a two-minute recording of a made-up language (Saffran, Aslin, & Newport, 1996). Evidence has even been found that sleeping newborns between 0.5 and 2 days old are able to detect statistical regularities in speech (Teinonen, Fellman, Näätänen, Alku, & Huotilainen, 2009).

Switching modalities we investigate visual sequential learning, a form of statistical learning that revolves around the recognition of associations and rules governing the presentation of a particular sequence. For instance, given a sequence where a blue triangle will always follow a yellow square, with enough exposure to the sequence we would be able to predict that given a yellow square the next shape in the sequence would be a blue triangle. In a sequence formed solely by rules such as these, it was found that 2-8 month olds can detect violations in the sequence (Kirkham, Slemmer, & Johnson, 2002) Evidence of the application of this type of visual sequential learning has been found in newborns as early as 1-3 days old (Bulf, Johnson, & Valenza, 2011), suggesting that visual, in addition to auditory, statistical learning may be functional from birth.

However, while it is clear that associative statistical learning is present from a very young age, it is not clear whether an understanding of temporal order within sequences is conserved. In other words, are these same subjects able to understand that the yellow square temporally precedes the blue triangle or have they simply made an association between the yellow square and the blue triangle, indifferent to the temporal relationship between the two? If so, are they thus liable to predict that the yellow square as being equally as likely to come after the blue triangle?

It has been shown that much like how 3-dimensional space is represented as ‘place cells’ and ‘grid cells’ (Hafting, Fyhn, Molden, M. Moser & E. Moser 2005; O’Keefe and Dostrovsky 1971), visual stimuli appear to have their own type of representational system within our neural structures. This is to say that seeing the yellow square causes one subset of neurons to fire while seeing the blue triangle causes a second subset of neurons to fire. Based on recording studies involving the perirhinal cortex [PRC] of primates we have come to understand that the PRC is critical for making associations between visual stimuli (Erickson & Desimone 1999; Miyashita 1988) where it was found that subsets of neurons fire more similarly when two stimuli were closely associated.

Recently, Schapiro, Kustner and Turk-Browne (2012) have shown that complex visual objects like fractals also hold similar neural representations within the hippocampus such that the neural representations of associated fractals become more and more similar to each other after exposure to visual sequential regularities. Human adults were placed in an fMRI and ‘actively viewed’ a 40-minute sequence of fractals while they were given the task of finding grey-scale patches randomly placed in a number of fractals. Unbeknownst to the subjects, the fractal sequence followed a set of rules where certain fractal pairs were highly associated, some weakly associated and some were not associated. Schapiro et al., found that the stronger the association between two fractals, the more similar their neural responses (recording of 81-191 voxels analyzed to 27-voxel cubic searchlight). Further, while every other region of the medial temporal lobe showed symmetrical changes for neural representations of associated fractals, the hippocampus (specifically the CA2/CA3/DG [trisynaptic circuit]) revealed asymmetrical changes. That is to say that while in most regions the initially distinct neural representation of the yellow square and blue triangle would meet a halfway point where they would eventually come to share very similar neural representations, the hippocampus was found to be different in that the neural representation of the first fractal would in fact come to take on the neural representation of the associated second fractal in a manner that was much quicker than the second fractal would come to take on the neural representation of the preceding first fractal. This phenomenon is particularly interesting because it is currently understood that the CA3 of the trisynaptic circuit is responsible for 'predictive firing' of place cells when mice navigate mazes they had previously learned (Johnson & Redish, 2007). Based on such findings, we believe that the trisynaptic circuit of the hippocampus is potentially responsible for encoding the temporal order of patterned-stimuli and may support a connection between 'forward-prediction' and the associative 'forward-order' for spatial and temporal tasks.

At the cellular level, however, the CA3 of rhesus macaque is not fully developed until after 1-year of age. This could potentially mean that the human CA3 does not fully develop until 3-4 years of age (Jabès, Banta-Lavenex, Amaral, & Lavenex, 2011). At a behavioral level, human children under 43-months are unable to complete a complex spatial memory task which involved memorizing the locations of several rewards with respect to spatial cues (Ribordy, 2013). Assuming shared processing between the neural representations of 3-dimensional space and those that manage temporal order, we hypothesize that associative sequential learning also follows a similar timeline for the development of the circuitry responsible for spatial processing. Based on this, we believe that children under the age of 43 months [immature-hippocampal] will be able to perceive that two shapes are associated with one another but they will not be able to discriminate details pertaining to temporal order. On the other hand, we hypothesize that children 43-months and older [mature hippocampal] will be able to not only understand associations but also perceive different rules of temporal order. This study will initially investigate the adult ability to perceive temporal order.

In order to begin our investigation into the effect of differential hippocampal development on sequence learning and pattern prediction, we opted to first pilot our paradigm on adults. Our expectation was that by the ages of 18-22, the hippocampus should be well outside of its developmental timeline (Jabès et al., 2011) and thus we would expect adults to successfully discriminate between forward and backward-order associations. After creating a successful experiment for adults we will then make small adjustments to our paradigm to suit testing in younger age groups.

General Methods for Experiments 1 and 2

For our adult paradigm, we chose an array of 6 abstract, nonsense shapes modeled after shapes used in previous visual statistical learning studies (Fiser & Aslin, 2001) (Figure 1). We opted to use black shapes presented against a white background in order to avoid age-specific learning techniques that older children/adults could utilize (i.e. learning the pattern through the colors of the shapes). The 6 shapes were pseudo-randomly divided into 3 pairs that were to be considered the 'Forward-Order Strong' [FOS] pairs (Figure 2A: when the first shape in the pair [A-shape] appears, the second shape in the pair [B-shape] always appears immediately afterwards, transitional probability: 1.0). When the FOS pairs are placed in a sequence, the second shape in a pair will then lead into the first shape in a second pair. This pair-within-pairs was deemed the 'Forward-Order Weak' [FOW] pair (transitional probability: .50) (Figure 2B). Note that nowhere in the sequence would a 'Backward-Order Strong' [BOS] pair (transitional probability: 0) naturally occur (Figure 2C). The pairs were then placed into a pseudo-randomly ordered sequence so as to avoid immediate repeats of pairs (e.g. A1B1A1B1). Figure 3 depicts a 9-shape long sample sequence. The pairs were balanced in a way so the subjects would be given equal exposure to each individual shape as well as each pair. The sequence would always begin on the second shape in a strong pair and then end with the first shape in a strong pair. This was done to avoid early segmentation of pairs.

Prior to beginning the training phase the instructions for the training task was read aloud by the experimenter to the subject. The subject was asked to give a verbal indication of their understanding before proceeding to a practice set of trials with standard shapes (i.e. triangle, circle, square, etc.), shapes that were not a part of the real training. After completing the practice trials the subject was allowed to ask any questions before beginning the actual training. Prior to beginning the training phase, the experimenter left the experiment room and the subject completed the remainder of the experiment alone. Subjects were not told they would be learning a sequence and were only told to complete the cover task of the given training task.

Upon completion of the experiment, the subjects of both learning tasks were asked a series of questions. One particular question involved laminated cards with the shapes from the experiment and the subject was asked whether they could recall any pattern given the shapes from the experiment. 2 of these 20 subjects were able to recall the exact pairs (1 subject from each learning task) while 4 of the 20 were able to partially recall the patterns (all from the Search-and-Find task).

Figure 1: The six abstract, nonsense shapes utilized in this study.

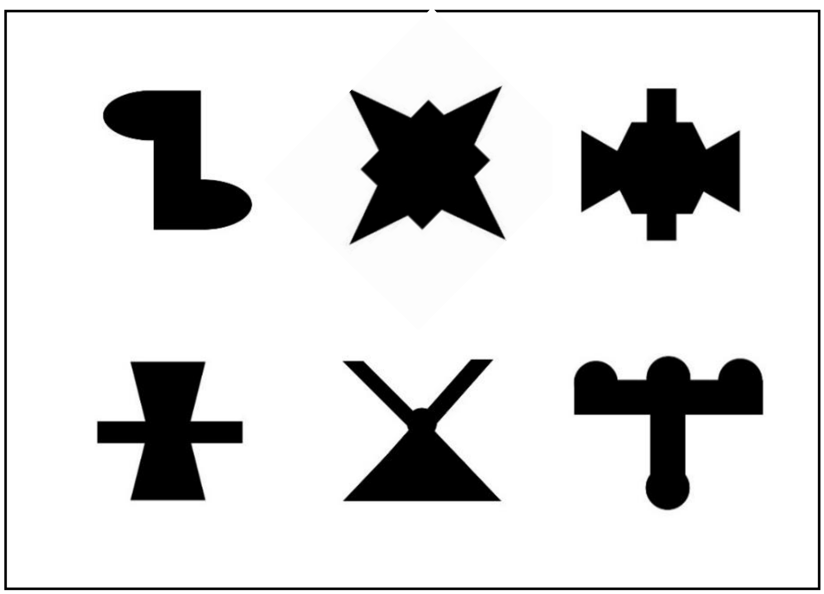

Figure 2: The different shape transitions.
Transitional probability of FOS= 1.0, FOW = 0.50, BOS = 0.

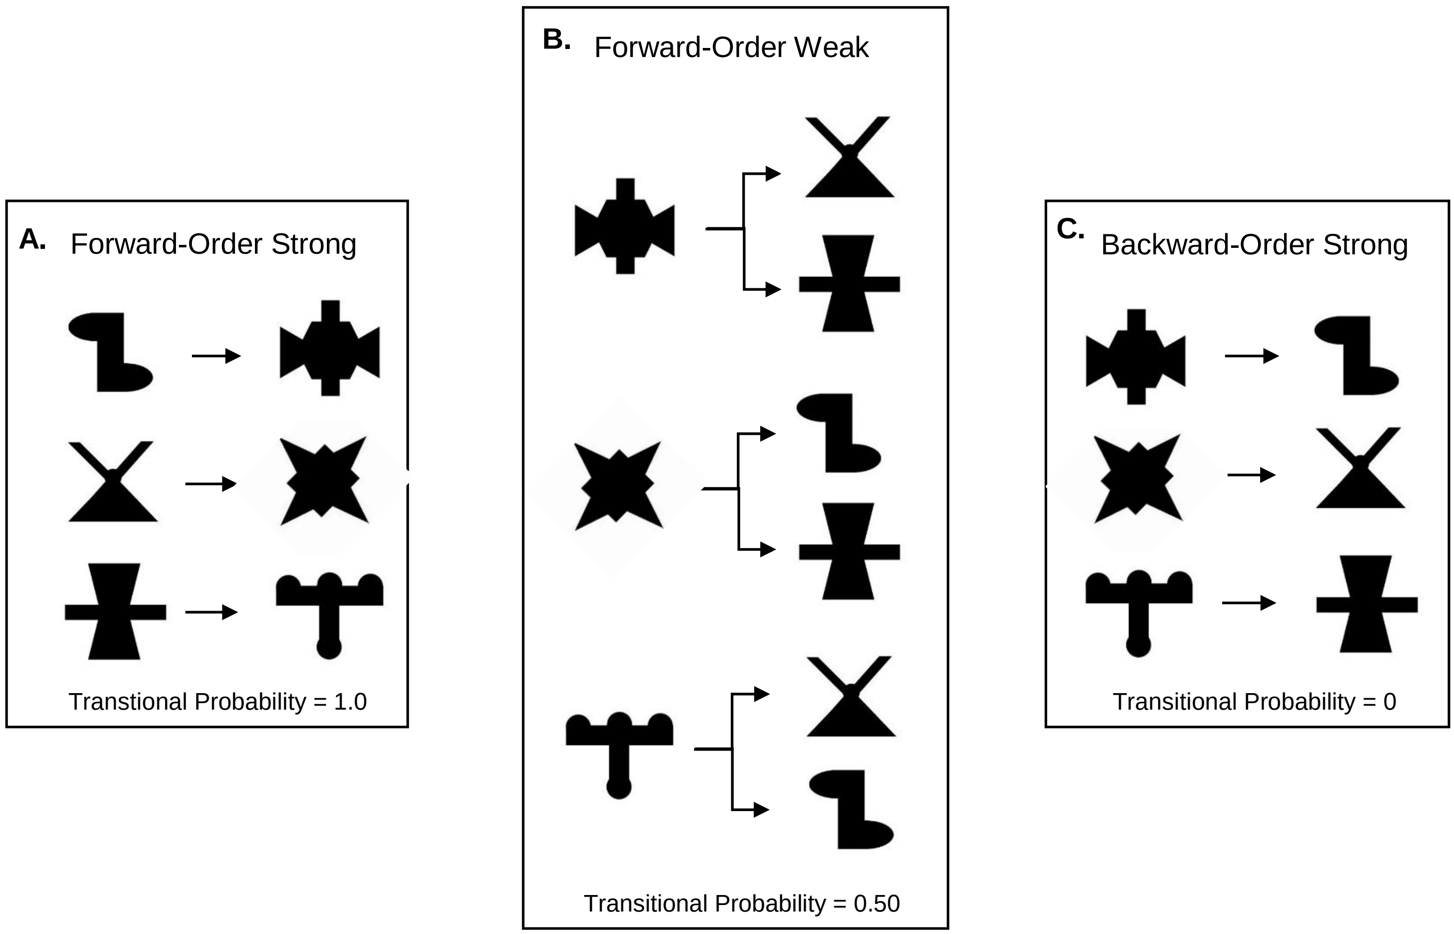

Figure 3: Depiction of how a sequence would be constructed. Notice no possibility of a backward-order transition.

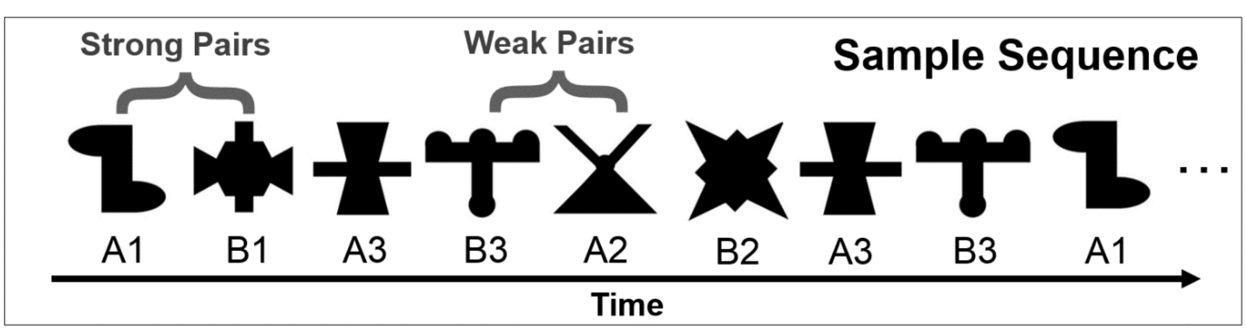

Experiment 1

Participants

Ten adults participated in Experiment 1. The participants were undergraduate students at the University of Arizona (9 females; range = 18-22 years). All students were enrolled in an introductory Psychology course and were participating for course credit. Each subject self-registered to participate in the study. Informed consent was obtained from all participants.

Materials

The six abstract shapes were created using Microsoft Word. These images were displayed on a Viewsonic VX1932wm display at 1440x900 resolution. The experimental paradigm was created and run using E-Prime 2.0.

Procedure

Training Phase: *Fade-to-Reveal*

The first rendition of our paradigm was termed ‘Fade-to-Reveal’ (Figure 4). It involved an exposure session where the subject was given a random array of all six shapes displayed on the screen at once. The shapes were initially presented ‘faded’ (grey) and each occupied 9x9cm region of space. The shapes were placed in one of six equally divided portions of the screen and their positions were pseudo-randomly selected (no shape was allowed to fall in the same place on consecutive trials). The subject would first attend to the six grey, faded shapes displayed against a white background and after 250 milliseconds, one shape would be ‘revealed’ by turning black. The subject’s task was to use the mouse to click on the revealed shape as quickly as they could. Pressing on the other five grey shapes did not allow for the trial to progress. 750 milliseconds after pressing on the black shape the trial would end and a new assortment of the six faded shapes would appear on the screen to which after another 250 milliseconds, the next shape would be revealed. This cycle would repeat until the subject completed the full sequence of shapes. Every subject of the Fade-to-Reveal paradigm was exposed to a sequence that was 504-shapes long. Within this 504-shape sequence, each of the six shapes were revealed a total of 84 times. The order in which the shapes were revealed corresponded to the sequential order of the pairs. For each trial we recorded how long the subject took to press on the revealed shape, the timer began immediately after the shape was revealed and the timer stopped once the subject pressed on the shape.

Figure 4: Schematic depicting a cycle of the Fade-to-Reveal task.

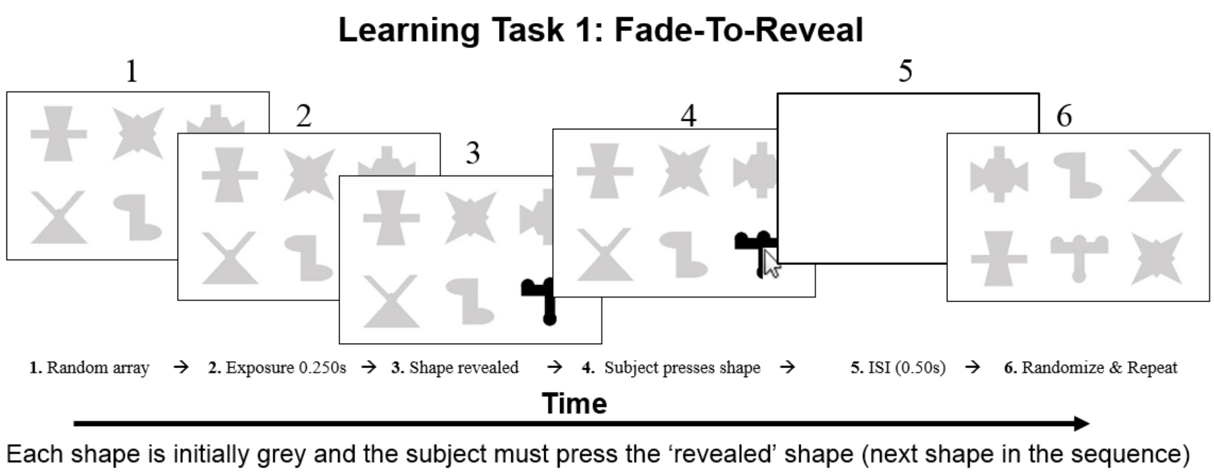

Test phase

After the training paradigm, the subjects were tested on their ability to predict the second shape in a pair given a preceding shape (Figure 5). The test paradigm was very similar to the learning paradigm in that the first screen of the trial showed a random array of the 6 grey shapes. One shape would be revealed by turning black and the subject was tasked with pressing the revealed shapes until they pressed on the target shape of that particular trial. That shape would turn black for 750 milliseconds and then a white screen would appear with the words 'Which shape comes next?' displayed in the center. The 6 shapes would then appear in black in the same array as they were when they were grey. The subject was instructed to select the shape they believed to be the next in the sequence. If the subject found it difficult to select a shape, they were told to make their best guess. After selecting a shape there was an inter-stimulus interval for 500 milliseconds (white screen) and then a new array of the six grey shapes was shown on the screen. Each shape was tested 4 times for 24 total test trials. The shapes were given in a random order that did not necessarily follow the pairings found in training. The only limitation placed in the presentation of the tested shapes was that no one shape was tested consecutively. During this we recorded which shape the subject predicted would come next as well as how long it took them to select their predicted shape. Subjects could complete a pair to make a Forward-Order Strong response (FOS), Backward-Order Strong (BOS) or Forward-Order Weak (FOW) with FOS pairs mapping onto the correctly predicted B-shapes following an A-shape, BOS mapping onto an incorrectly predicted A-shape associated with (but normally preceding) a B-shape, and FOW one of two correctly predicted A-shapes following a B-shape.

Figure 5: Schematic depicting a cycle of the test given to subjects after the learning phase.

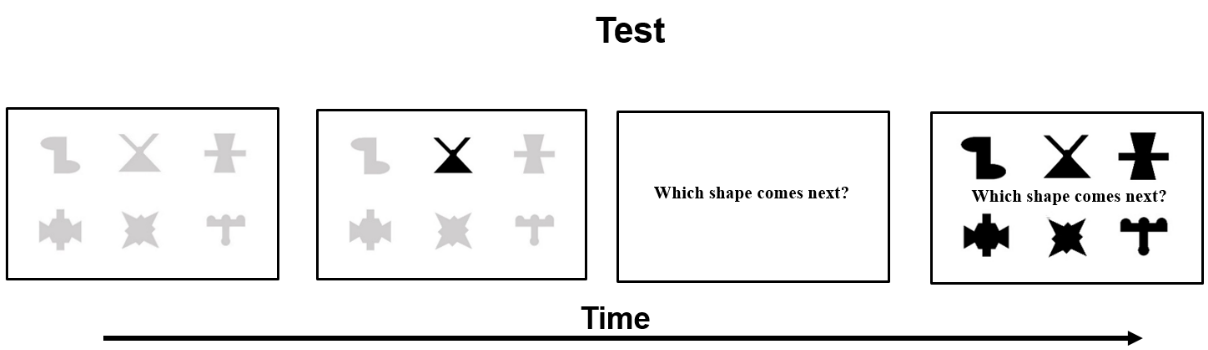

Results

Training Phase: *Fade-to-Reveal*

Reaction time [RT] to press the revealed shape was used as the dependent variable indicating performance through the trial. A repeated measure ANOVA with 11 within-subject levels (for training block) was performed to observe performance for the A-shapes compared to the B-shapes. Analysis of performance across the 11 training blocks for A-shapes revealed significant changes in RT across training, *F*(10,90)=2.957, *p*=003. Polynomial trends analysis revealed significant order 5, 8, and 10 trends, *F*(1,9)=*9.524, p*=.013, *F*(1,9)=6.263*, p*=.034, *F*(1,9)=13.792*, p*=.005 respectively, suggesting performance for A-shapes followed a sawtooth-pattern (Figure 6A). Analysis of performance of the B-shapes across the 11 training blocks also revealed a significant effect of set, *F*(10,90)=2.720, *p*=.005. Polynomial trends analyses for the B-shapes reflected significant cubic trends indicating two inflection points and thus two changes in direction for performance through training. (Set 4 and Set 11), *F*(1,9)=31.332, *p*=.000 for the cubic trend (Figure 6B)

Figure 6: Graphs depicting reaction time through training. The 504-shape long training phase was broken down into sets. Each set contained 12 shapes (6 A-shapes and 6 B-shapes) with sets spaced approximately 50 shapes apart (Set 1: Shapes 1-12, Set 2: 50-61, Set 3: 100-111, Set 4: 150-161, Set 5: 200-211, etc.) Due to an uneven number of shapes, the final set (Set 11) spanned from shape 493-504.

A. Performance for A-shapes depicted across 11 training sets.

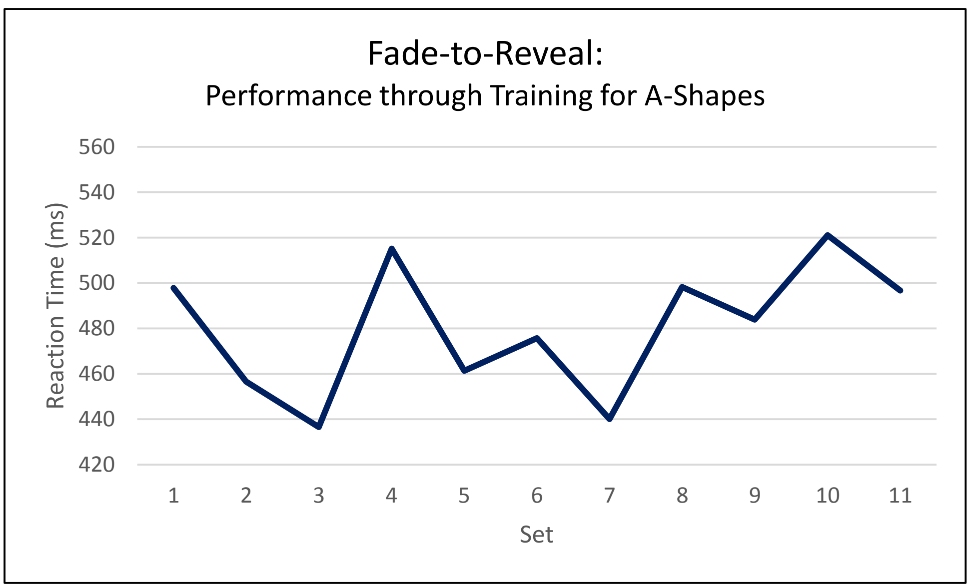

B. Performance for B-shapes depicted across 11 training sets

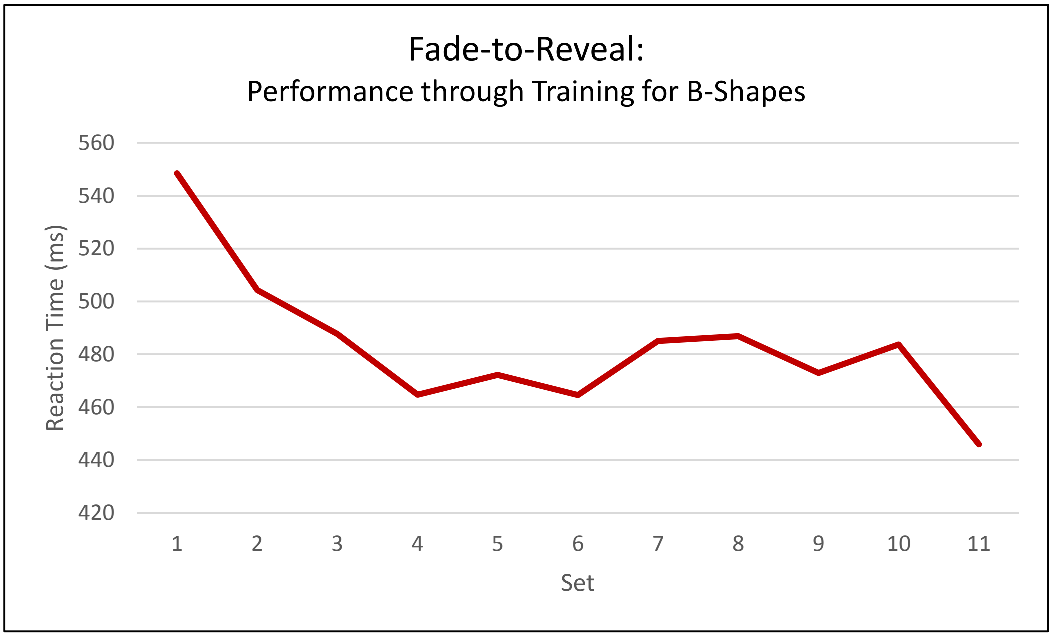

C. Composite of A-shape performance and B-shape performance across training.

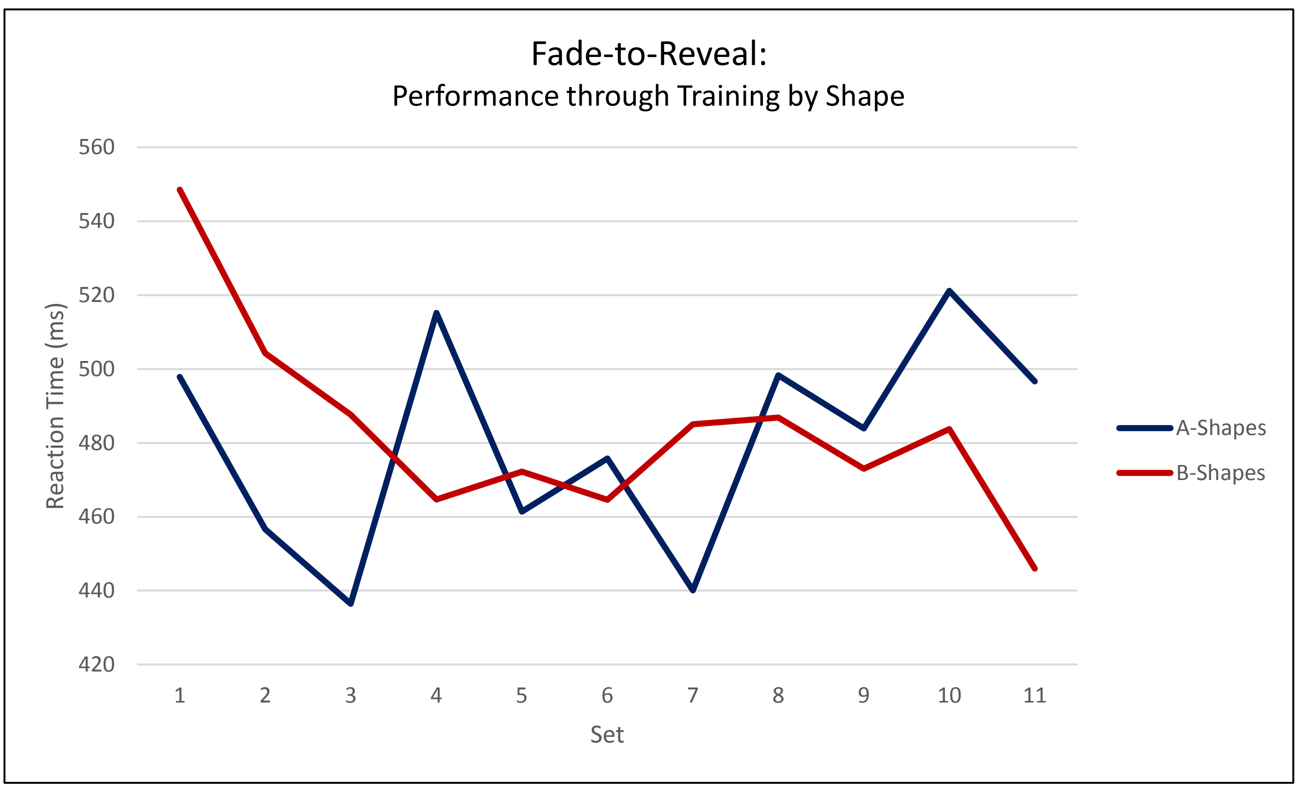

Test Phase

We conducted a paired t-tests comparing the percentage of predictions for FOS to BOS and FOS to FOW (Figure 7). There was no difference in mean accuracy for FOS vs. BOS, $t(9)=.90$, $p=.29$**,** nor was there a difference for FOS vs. FOW, $t(9)=.90$, $p=.39$.

Figure 7: Graphical depiction of the order-recall test after the Fade-to-Reveal training task.

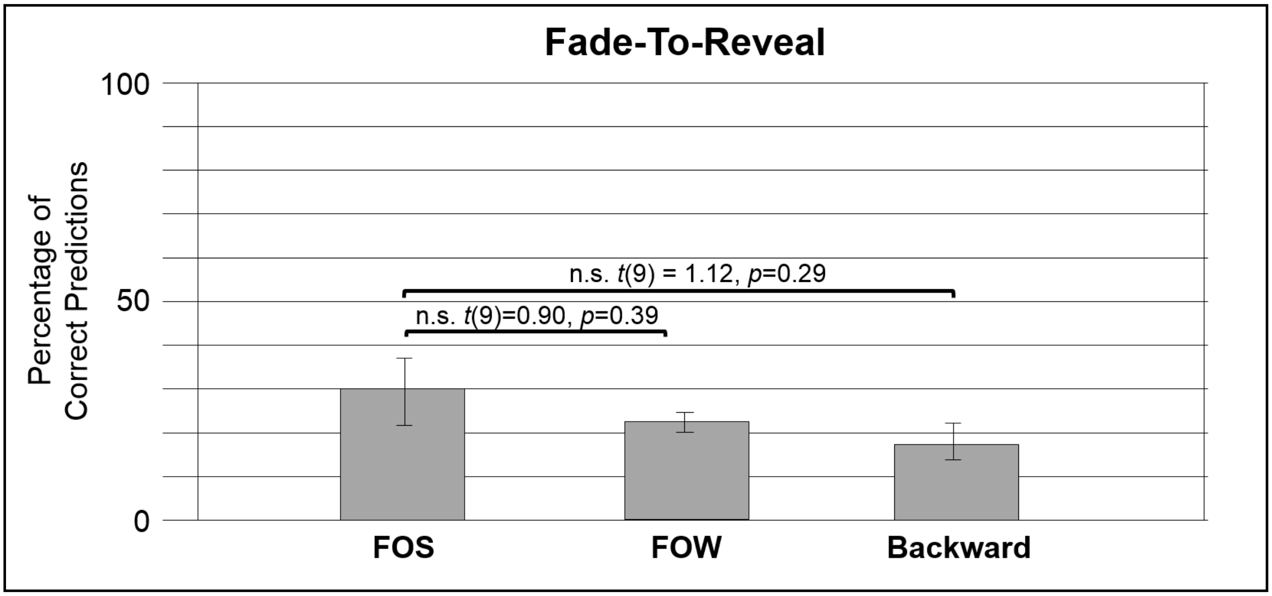

Based on the results of this test, the ten subjects of this training task were then ordered, least to greatest, by their number of ‘No-order’ predictions (a prediction was labeled ‘No-order’ when it did not fall into the three other orders, i.e. given shape A1, the subject responds with A2. A1 -> A2 does not fall into a FOS, FOW or BOS order). In the case of equal number of No-order predictions, a secondary measure of least number of Backward-order predictions was used (Table 1). This group of ten subjects was then divided along the median with the group with the lower number of No-order and Backward (BOS) predictions being referred to as the ‘High-performers’ and with the group with the higher number of No-order and BOS predictions being ‘Low-performers’. Based on these categories, a repeated measures ANOVA with 11 within-subject levels was performed for the two groups to compare reaction times for B-shapes across the 11 sets of training. The group of high-performers revealed significant improvement in RT across the 11 sets of training, $F(10,40)=2.868$, $p=.009$, following a linear and cubic effect, $F(10,40)=77.27$, $p=.001$ and $F(10,40)=18.960$, $p=0.12$ respectively (Figure 8A). The low-performers revealed no significant change in RT across training, $F(10,40)=1.193$, $p=.33$ (Figure 8B).

Table 1: High-performers vs. Low-performers from test. Five subjects with the least number of No-order pairings formed the group of 'High-Performers'. The remaining five subjects were placed in the group of 'Low-Performers'. Reaction time (ms) data for B-shapes in training was then extracted for the subjects. To the left of the subject number we find the number of No-order predictions with the number of Backward-order predictions in parentheses.

		Number of No-order Predictions (Number of Backward Pairings)	Subj.	Mean RTs by Performance										
				Set 1	Set 2	Set 3	Set 4	Set 5	Set 6	Set 7	Set 8	Set 9	Set 10	Set 11
B-Shapes	High Performers	0(1)	**102**	618.667	502.167	551.5	510.667	594.5	544.833	591.667	535.167	546.333	526.5	484
		12(1)	**104**	583.5	619.333	587.833	530.5	526.5	480.5	398.167	527.167	489.833	670.333	405
		12(3)	**106**	539	462	465.333	515.5	434.333	393.833	465.333	485.833	406	422.167	433.5
		13(3)	**103**	582.5	496	485.833	508.5	496.833	525	479.667	498	500.667	465.333	407.167
		13(6)	**105**	449.833	512.833	390.667	370.333	336.833	379.5	347.333	392.167	401.333	356.667	332.667
	Low Performers	14(3)	**107**	583.667	565.333	498.5	449.5	451.167	446	459.833	450.833	350.333	431.167	352.833
		15(1)	**109**	601.167	497.667	501.167	491.5	536	515.333	592.5	536.833	546.167	544.667	493.167
		16(2)	**101**	607.667	564.667	588.5	552.667	527	549.333	578.5	549.5	536.833	532.833	542.5
		17(0)	**108**	439.833	441.167	435.167	413	454	409.5	524.833	449.667	566.667	528.667	441.5
		17(2)	**110**	479.5	381.333	372.167	304.667	364.667	401.667	412.167	443.167	385.667	359	567.667

Figure 8: Graphs depicting performance on training for B-shapes based on performance on test.

A. Depiction of the average RT for B-shapes through 11 sets of training for those who performed well on the test.

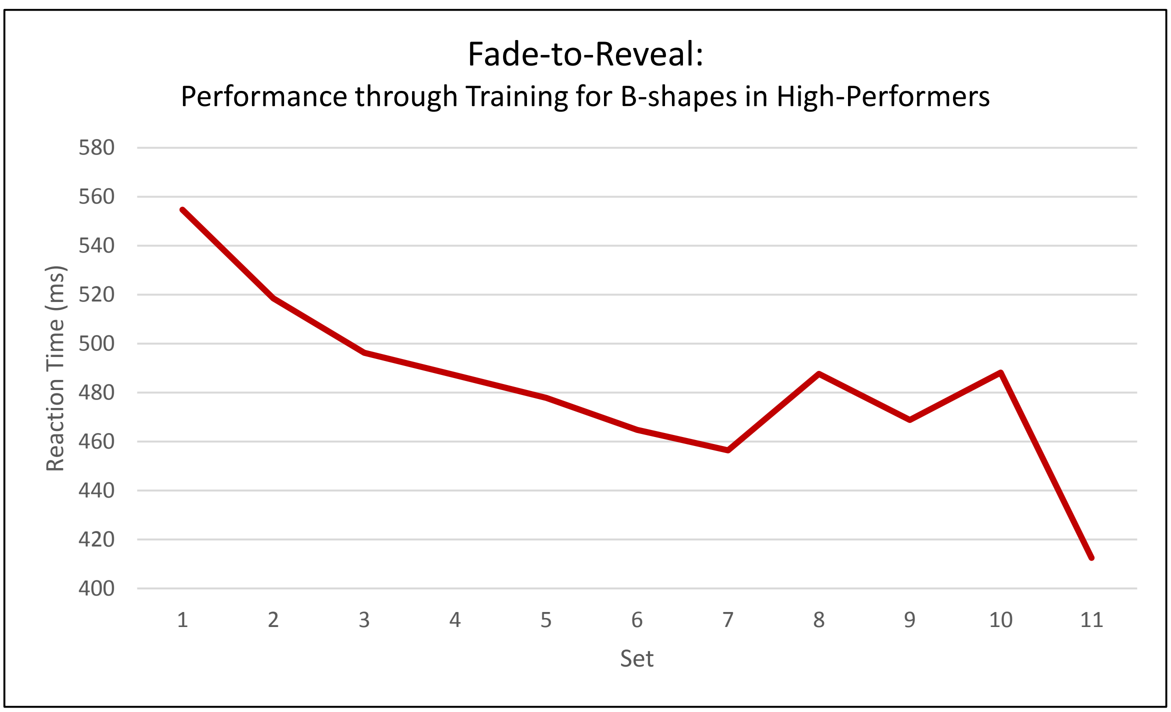

B. Depiction of the average RT for B-shapes through 11 sets of training for those who performed poorly on the test.

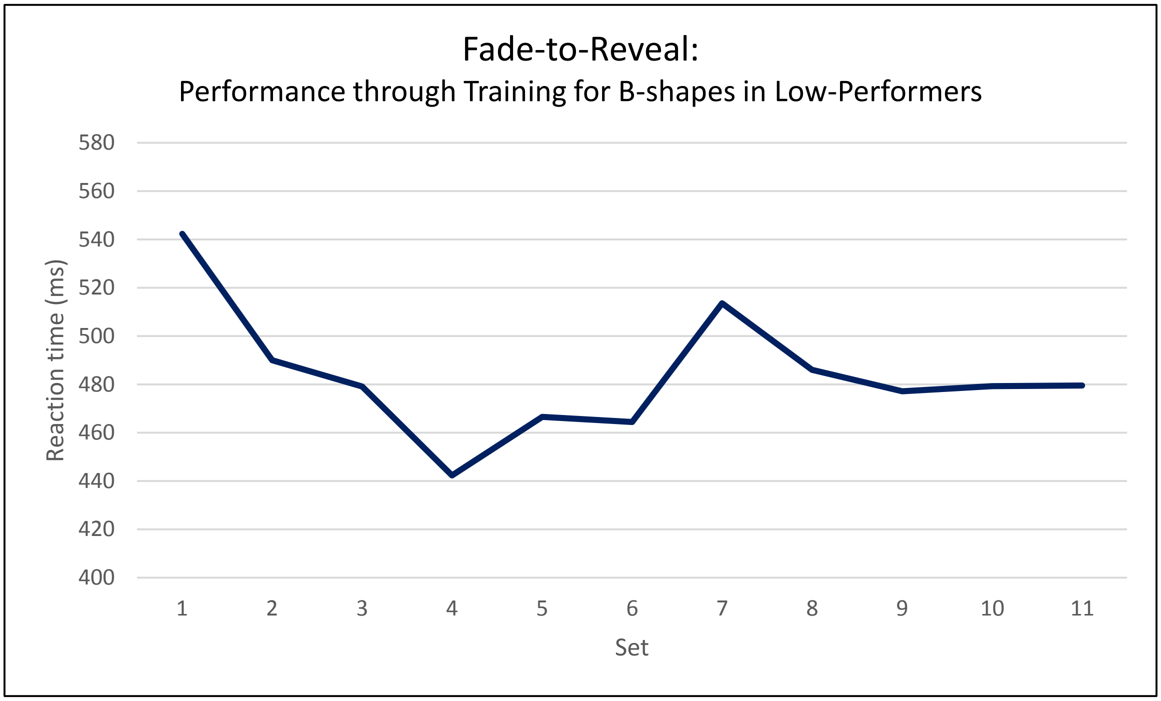

C. Composite of High-Performers and Low-Performers

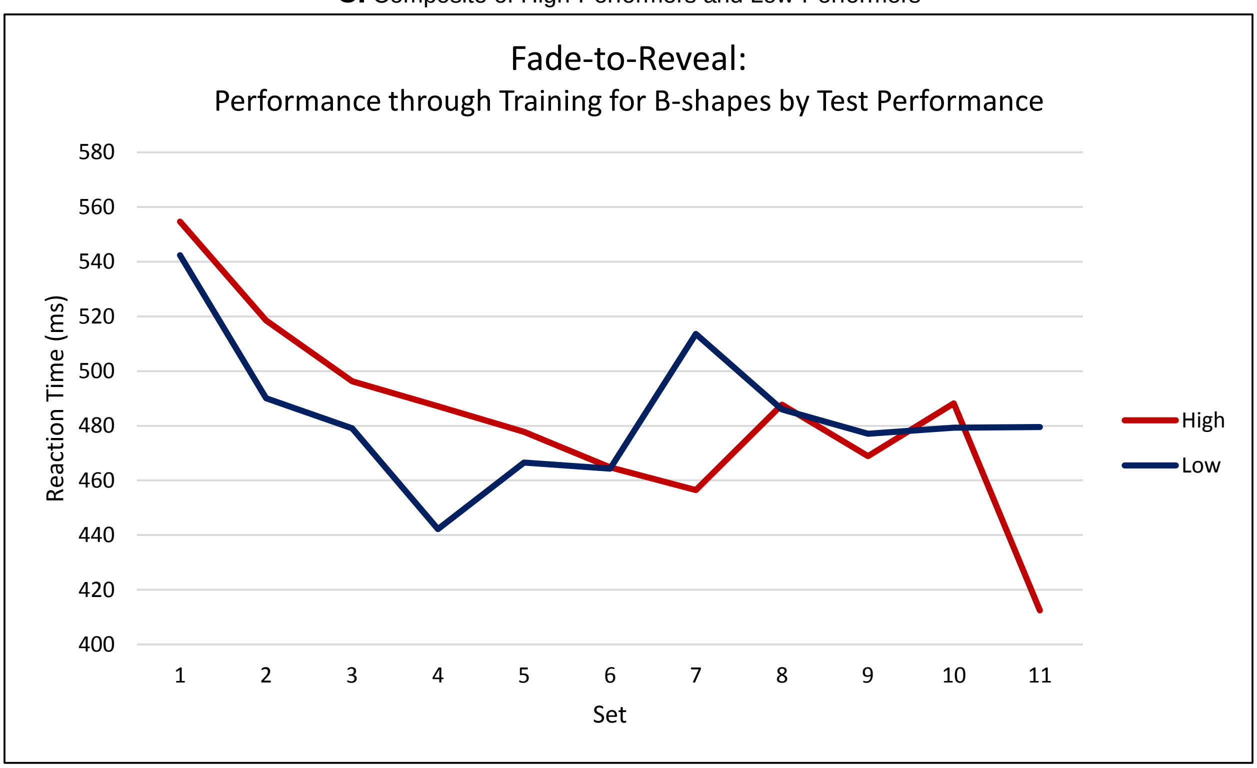

Discussion

In our analysis of performance for Shape A we observe a trend which reveals that performance for Shape A undulated throughout training. This rising and falling of performance suggests the subjects were not able to thoroughly learn the transition from Shape B to Shape A during training which is expected. In our analysis of performance for Shape B we observe two points of improved performance through training (Set 4 and Set 11). This cubic trend in the performance for Shape B suggests that the subjects were able to learn and consistently 'predict' the transition from an A-shape to a B-shape so as to show improvement throughout training.

This differential performance improvement in the two shapes is expected as learning and prediction for Shape A is more difficult considering the transitional probability from Shape B to Shape A is 0.50 (Forward-Order Weak transition). In comparison the transitional probability for Shape A to Shape B is 1.0 (Forward-Order Strong transition). Due to this considerable difference in transitional probability, we expected (and subsequently found) that subjects would be much quicker to learn and improve their performance for responding to Shape B than to Shape A. We believe these performance results reveal that learning for the B-shapes did in fact occur and the decrease in reaction time was not merely the product of the subject learning the paradigm. Had the improvement in B-shape performance been simply the result of the subjects becoming better at performing the learning task itself, the improvement would have been found in A-shape performance as well.

In the analyses for the results of the order-recall test, we find that when given one shape, subjects were unable to correctly predict which shape would proceed the given shape in the sequence. We find that not only did subjects have no real ability to discriminate between the three different orders but that subjects did not perform significantly above chance in their predictions of order. This appears to be in direct contradiction to the results of the training phase which suggests that, at the very least, learning for Forward-Order Strong pairs did occur. We believe that the results of the performance during the training phase in light of the results of the test indicates that the learning of the sequence during the training phase was implicit. When the subjects were given an explicit test, they were unable to utilize their implicit knowledge of the sequence to correctly predict which shape would proceed any particular given shape. This interpretation of these two results falls in line with our particular experimental paradigm as the subjects were not explicitly told that they would be learning a sequence. Instead, they were naive to the presence of a sequence and were told that their task during training was to press the revealed shapes as quickly as possible.

A second interpretation of these conflicting results in training and test could be that the results of the order-recall test were simply the product of a lack of power. This interpretation is reinforced by the analysis done comparing those subjects with lower No-order and BOS predictions in test against those who had higher No-order and BOS predictions. The group of those that performed better on the test also showed significant improvement in RT for B-shapes through training while on the other hand those that performed worse on the test showed no significant change in training. These results reveal a direct correlation between performance in training that would imply learning may not be implicit and that with more power we would see significance.

In an effort to understand if the Fade-to-Reveal task was simply too passive and led to low order-discrimination scores for this reason, we created a second paradigm. This new paradigm was similar to the Fade-to-Reveal task but involved more subject interaction. We predicted we would see more robust learning due to an increase in interaction with the shapes during training. We expected this to translate to higher performance on the test and more rapid improvement in B-shape performance during training.

Experiment 2

Participants

Ten adults participated in Experiment 2. The participants were undergraduate students at the University of Arizona (10 females; range = 18-22 years). All students were enrolled in an introductory Psychology course and were participating for course credit. Each subject self-registered to participate in the study. Informed consent was obtained from all participants.

Materials

The six abstract shapes used in this paradigm were the same shapes used in Experiment 1. The sequence the subjects were exposed to also remained the same. These images were displayed on the same Viewsonic VX1932wm display at the same 1440x900 resolution. The experimental paradigm was created and run using E-Prime 2.0.

Procedure

Training Phase: *Search-and-Find*

In our second rendition of the exposure paradigm, subjects were tasked with actively 'searching' for the target shape in an array of the six grey shapes (Figure 9). The act of searching consisted of using the computer mouse and clicking on the grey shapes until the subject pressed on the target shape of that particular trial. The target shape would turn black for 750 milliseconds to signify that it was 'found' and then the trial would end. The subject would see a white screen for 750 milliseconds before a new pseudo-random assortment of shapes appeared on the screen with which the subjects would repeat the 'Search-and-Find' task. The target shapes to be found corresponded with the shapes following a particular sequence of the 3 shape pairs. Subjects were not told they would be learning a sequence of shape pairs and were only instructed to complete the task of the given experiment to the best of their ability. For each trial we recorded the time it took for the subject to find the target shape and how many shapes they pressed before finding the target shape. All subjects of the Search-and-Find task were exposed to a 336-shape long sequence.

After the end of the training phase, the subjects were instructed that they would be moving onto Phase 2 of the experiment (the test phase).

Figure 9: Schematic depicting a cycle of the Search-and-Find task.

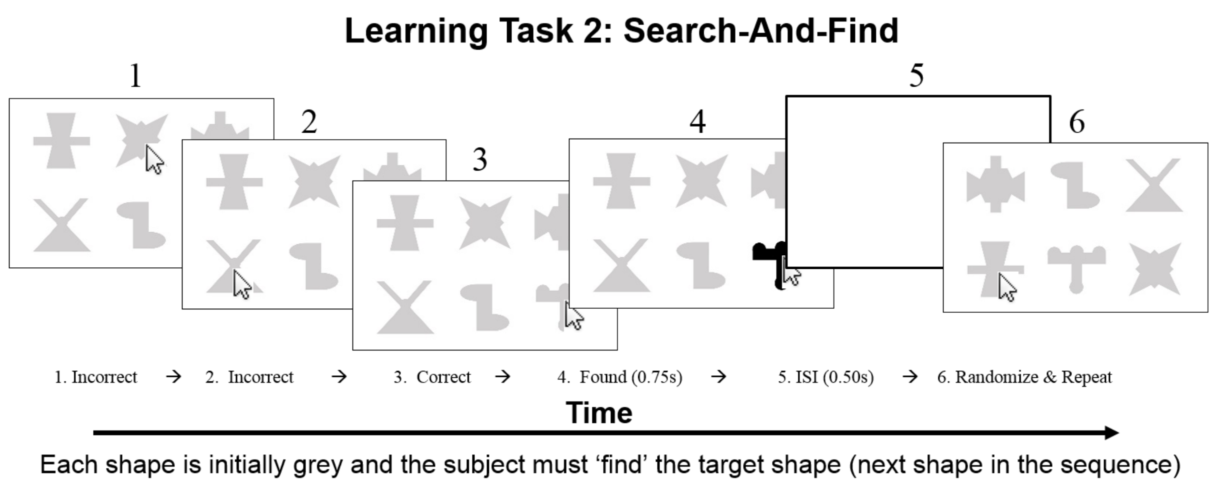

Test Phase:

After the training paradigm, the subjects were tested on their ability to predict the second shape in a pair given a preceding shape (Figure 5). The test paradigm was very similar to the learning paradigm in that the first screen of the trial showed a random array of the 6 grey shapes. The subject was tasked with pressing the different grey shapes until they pressed on the target shape of that particular trial. That shape would turn black for 750 milliseconds and then a white screen would appear with the words 'Which shape comes next?' displayed in the center. The 6 shapes would then appear in black in the same array as they were when they were grey. The subject was instructed to select the shape they believed to be the next in the sequence. If the subject found it difficult to select a shape, they were told to make their best guess. After selecting a shape there was an inter-stimulus interval for 500 milliseconds (white screen) and then a new array of the six grey shapes was shown on the screen. Each shape was tested 4 times for 24 total test trials. The shapes were given in a random order that did not necessarily follow the pairings found in training. The only limitation placed in the presentation of the tested shapes was that no one shape was tested consecutively. During this we recorded which shape the subject predicted would come next as well as how long it took them to select their predicted shape. Subjects could complete a pair to make a Forward Order Strong response (FOS), Backward Order Strong (BOS) or Forward Order Weak (FOW) response with FOS pairs mapping onto the correctly predicted B-shapes following an A-shape, BOS mapping onto an incorrectly predicted A-shape associated with (but normally preceding) a B-shape, and FOW one of two correctly predicted A-shape following a B-shape.

Results

Training Phase: *Search-and-Find*

The time necessary to find the target shape was used as the dependent variable. With a repeated measure ANOVA given 8 within-subject levels reflecting performance across the 8 training blocks, performance was analyzed separately for the first shape in a pair (A-shapes) and the second shape in a pair (B-Shapes). Across training with respect to A-shapes we found no significant changes in performance, $F(7,63)=1.234$, $p=.298$ (Figure 10A). With B-shapes, analysis across training for the 8 sets revealed a violation of sphericity given Mauchly's test, $\chi^2(27)=54.54$, $p=.003$. Due to this we used the Greenhouse-Geisser corrected test ($\varepsilon=.505$) and found that performance for B-shapes across training showed no significant change, $F(3.538,31.843)=1.538$, $p=.219$ (Figure 10B). Performance for A and B shapes is shown together in Figure 10C.

Figure 10: Graphs depicting reaction time through training. The 336-shape long training phase was broken down into sets. Each set contained 12 shapes (6 A-shapes and 6 B-shapes) with sets spaced approximately 50 shapes apart (Set 1: Shapes 1-12, Set 2: 50-61, Set 3: 100-111, Set 4: 150-161, Set 5: 200-211, etc.) Due to an uneven number of shapes, the final set (Set 8) spanned from shape 325-336.

A. Performance for A-shapes depicted across 8 training sets.

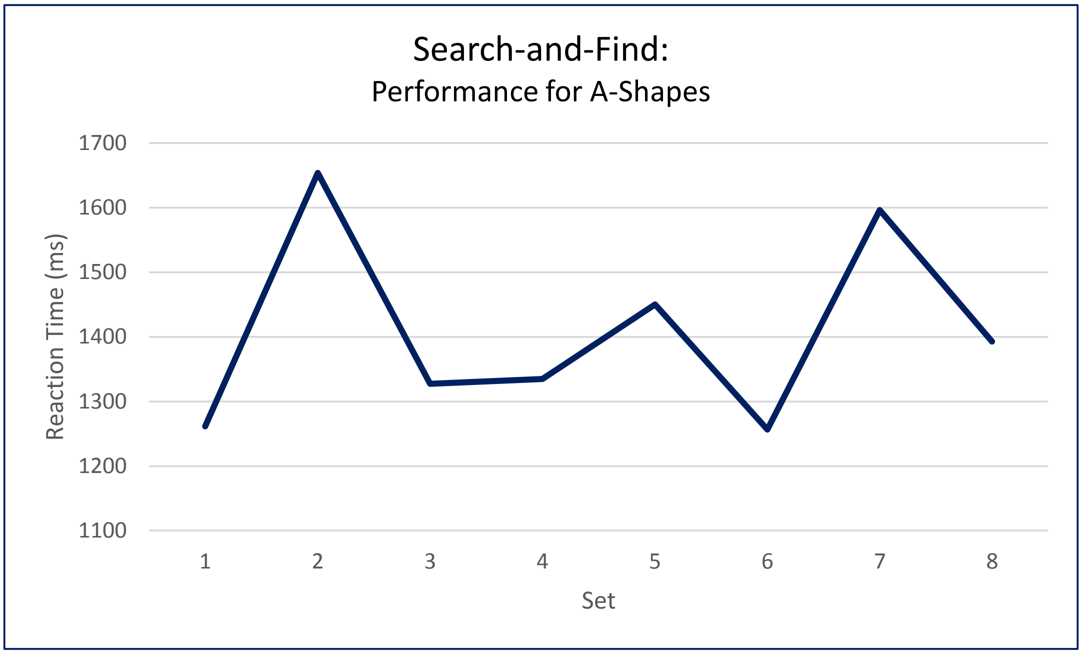

B. Performance for B-shapes depicted across 8 training sets.

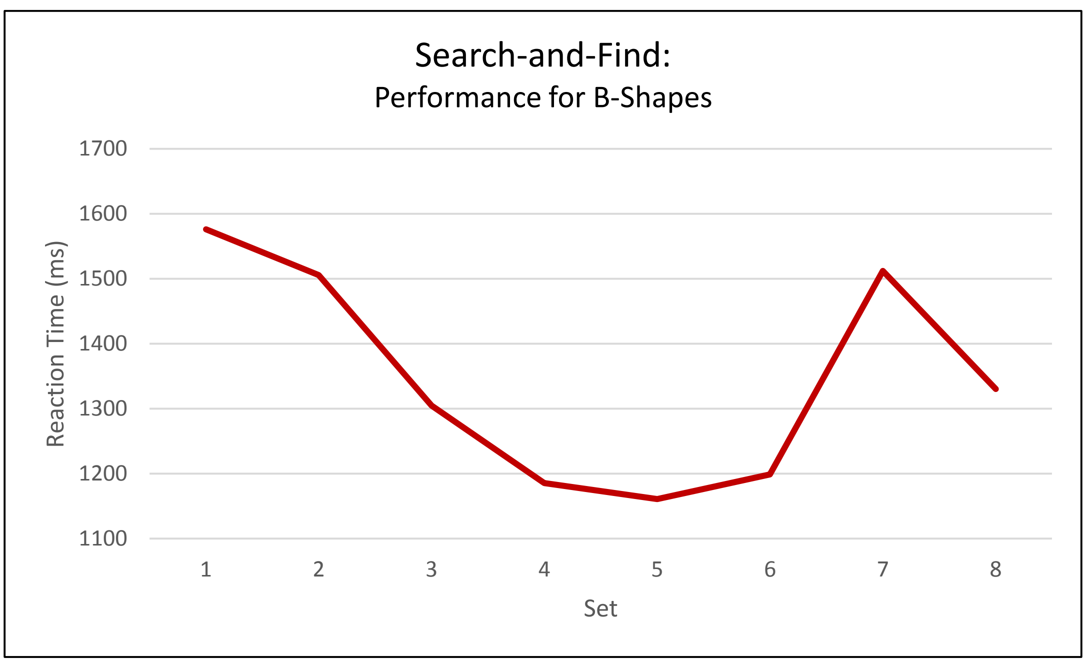

C. Composite of A-shape performance and B-shape performance across training

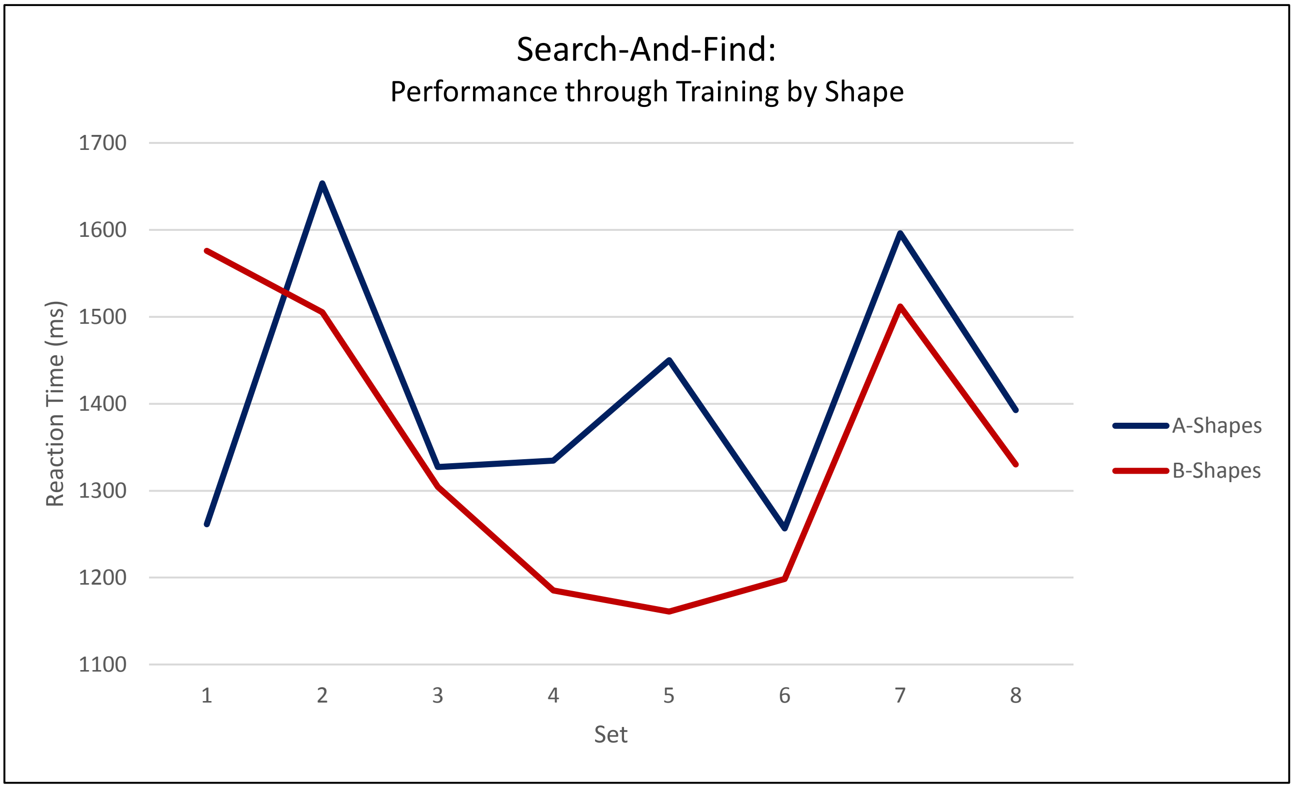

Accuracy

The number of incorrect presses was used as the dependent variable for determining the accuracy for a given trial. A repeated measures ANOVA with 8 within-subject levels reflecting performance across the 8 training blocks was conducted for both A-shapes and B-shapes. With respect to A-shapes we found no significant changes in the number of errors through training $F(7,63)=1.841$, $p=.095$, (Figure 11A). Analysis of the number of errors for B-shapes through training reveals a significant quadratic trend, $F(7,63)=3.161$, $p=.006$ and $F(1,9)=1.968$, $p=.001$ for the quadratic trend, indicating a single inflection point and thus one change in direction through training reflecting the fact that errors initially decreased then increased during training (Figure 11B). Accuracy for A and B shapes is shown together in Figure 11C.

Figure 11: Graphs depicting the number of errors in the particular set across training. Every mouse click on a shape that was not the target shape of that particular trial was counted as an error.

A. Errors for A-shapes across the 8 training sets

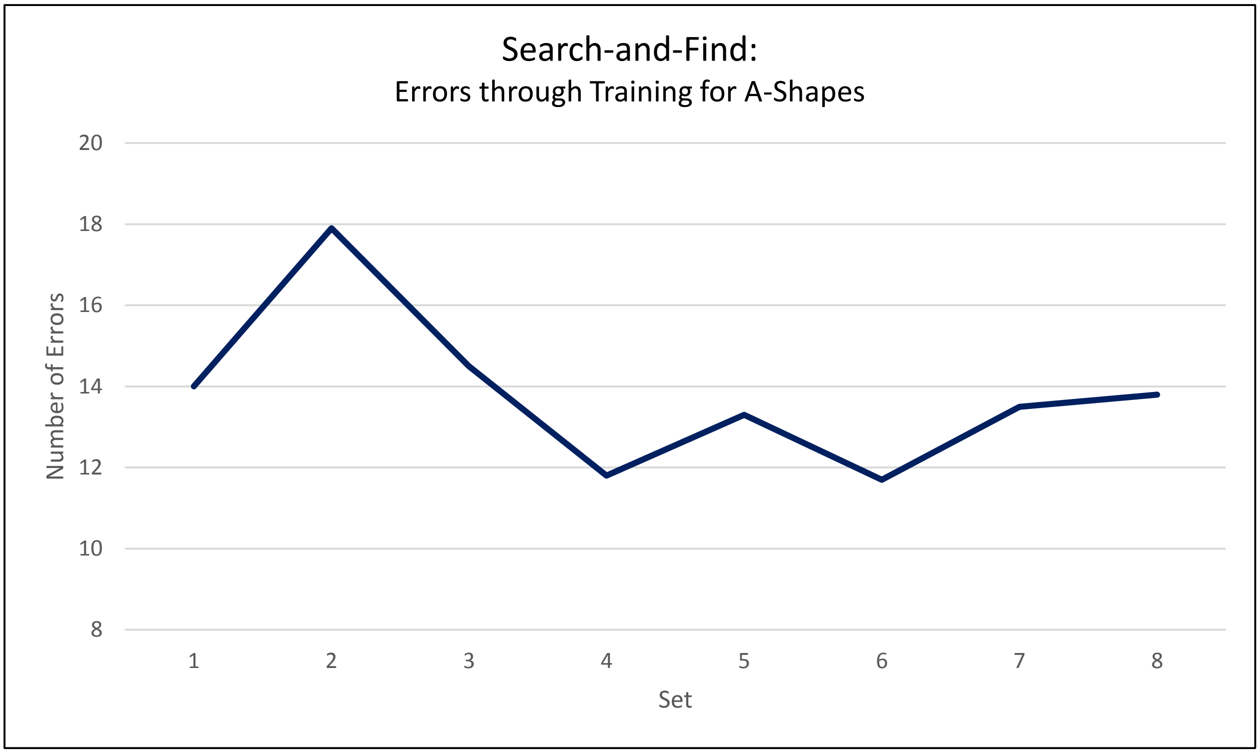

B. Errors for B-shapes across the 8 training sets

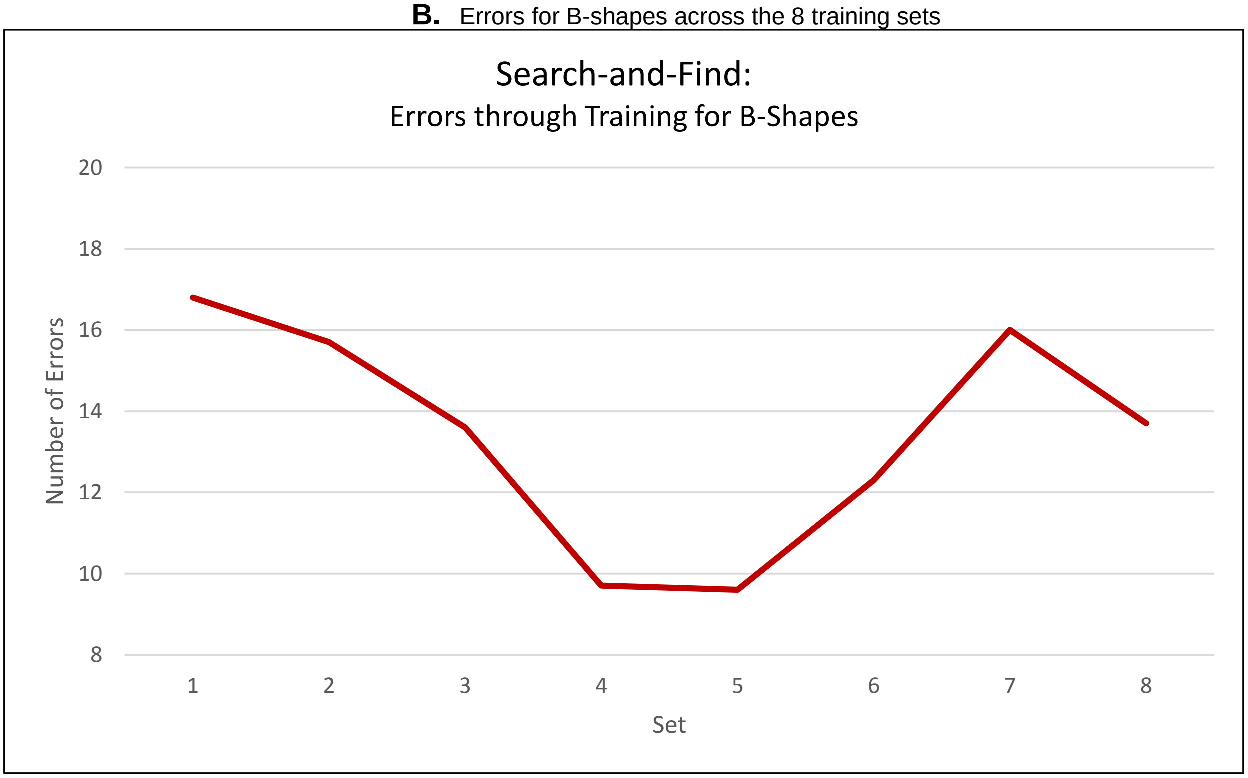

C. Composite of A-shape errors and B-shape errors overlapped with performance for both shapes. Errors shown in dashed lines

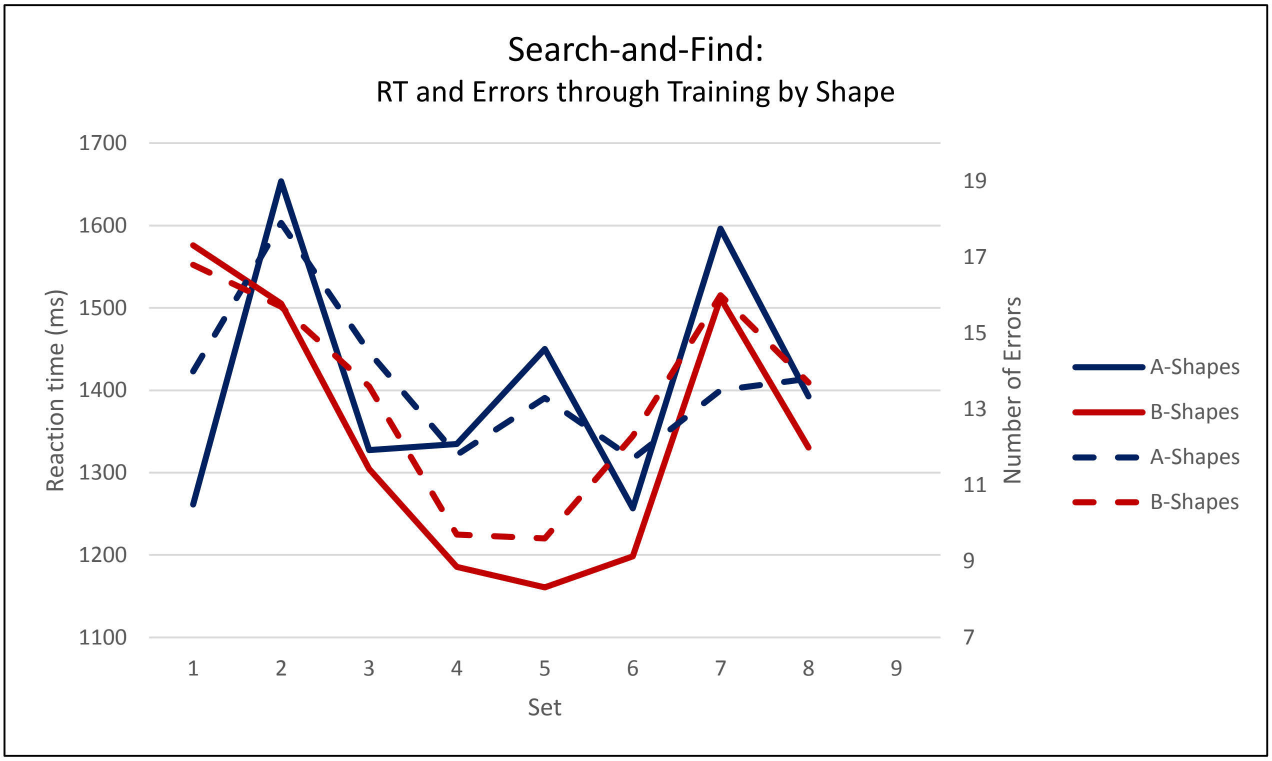

Test Phase

We conducted paired t-tests comparing the percentage of predictions for the three different orders. We compared predictions of FOS to BOS and FOS to FOW (Figure 12). There was no significant difference in correct predictions of FOS vs. BOS, $t(9)=.845$, $p=.43$, nor was there a difference found in the comparison of FOS predictions to FOW predictions, $t(9)=1.461$, $p=.18$.

Figure 12: Graphical depiction of the order-recall test after the Search-and-Find training task.

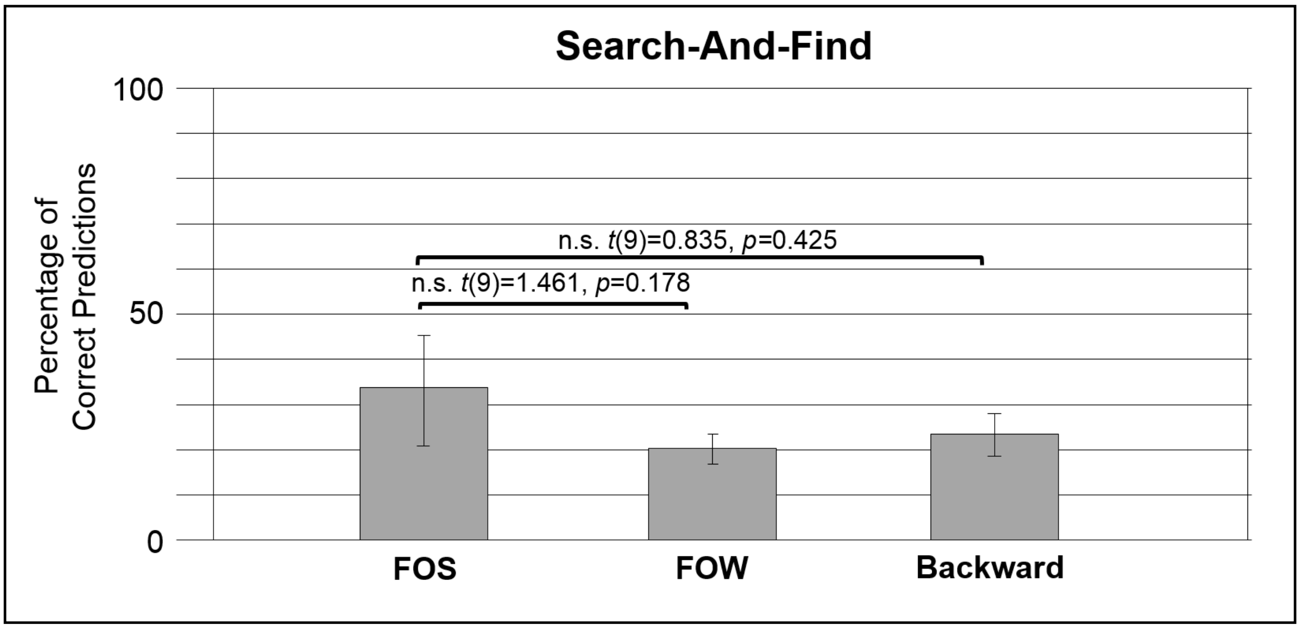

As was done in the Fade-to-Reveal results, we separated the subjects by the number of No-order and Backwards-order predictions. Based on these categories, a repeated measures ANOVA with 8 within-subject levels was performed for the two groups to compare reaction times for B-shapes across the 8 sets of training. The group of high-performers revealed no significant change in RT across the 8 sets of training, $F(7,28)=1.487$, $p=.212$ (Figure 13A). The low-performers revealed no significant change in RT across training as well, $F(7,28)=1.551$, $p=.191$ (Figure 13B). Performance on A and B shapes is shown together in Figure 13C.

Figure 13: Graphs depicting performance on training for B-shapes based on performance on test.

A. Graphical depiction of the average reaction time for B-shapes in training for those who performed well on the test.

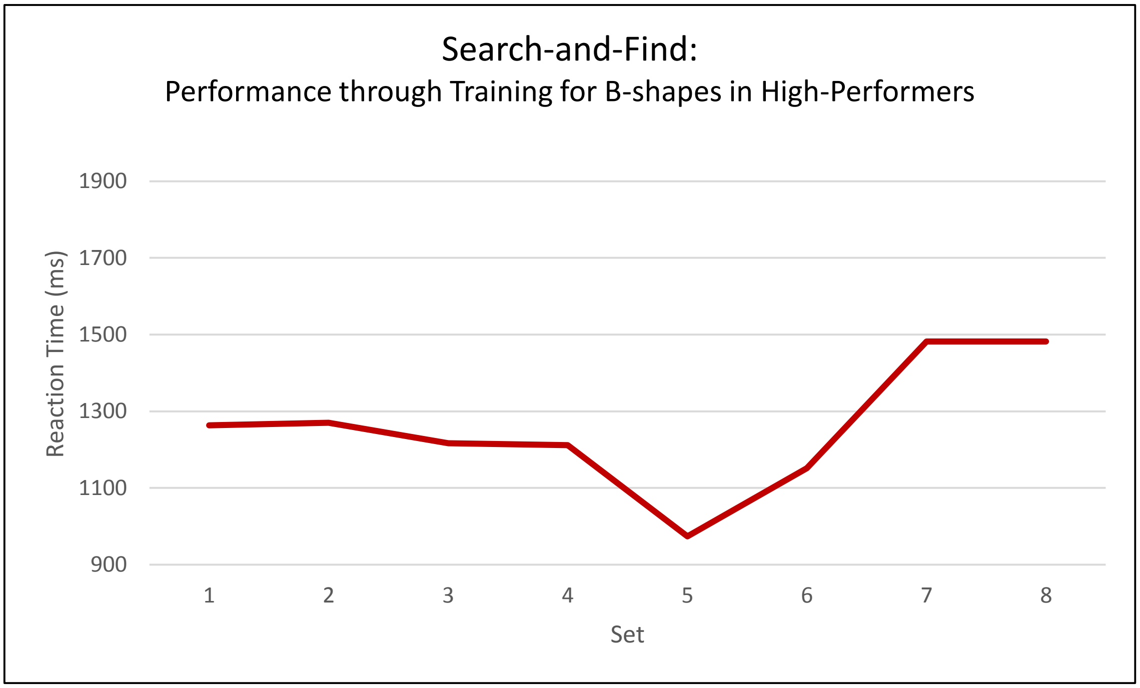

B. Graphical depiction of the average reaction time for B-shapes in training for those who performed poorer on the test.

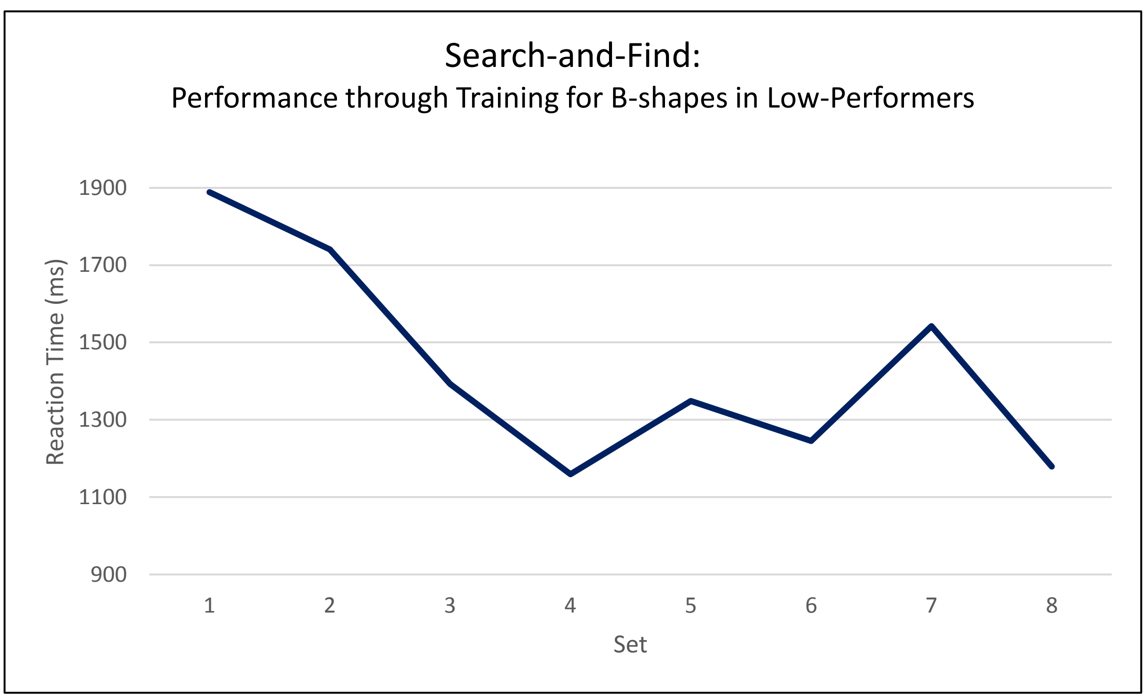

C. Composite of High-Performers and Low-Performers

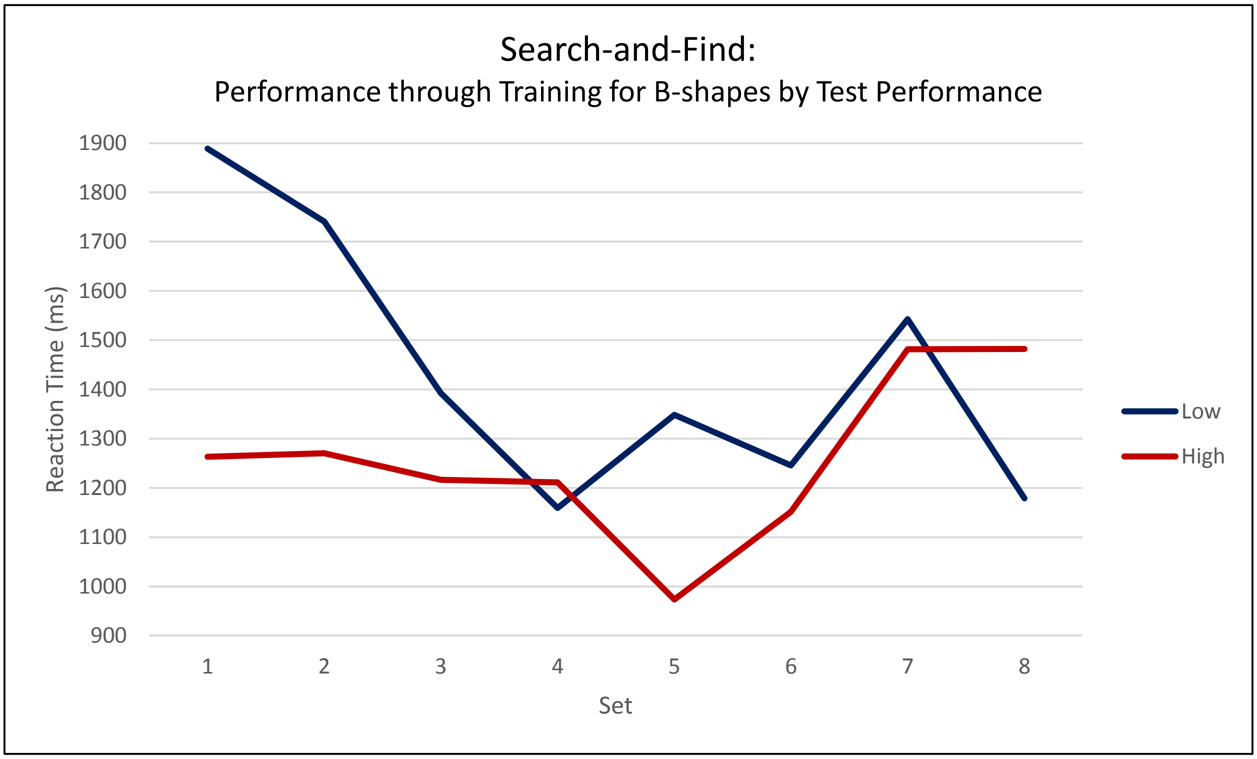

Discussion

While performance on both A-shapes and B-shapes across training showed no significant changes in performance, in an effort to explore the data we investigated for the performance for B-shapes through polynomial trend analysis. Performance for B-shapes was found to follow a quadratic trend, $F(1,9)=8.124$, $p=.019$. This reveals that reaction times for B-shapes during training initially decreased but then increased. While ultimately the results of B-shape performance were insignificant, we interpret this quadratic effect of training to indicate that subjects initially showed an ability to at least learn the paradigm but then the subsequent increase in reaction time suggests fatigue. Had the subjects learned the sequence we would expect that, with the onset of fatigue, the error rate would remain low but the time needed to find the target B-shape would increase. It is possible, however, that as fatigue set in subjects employed a different strategy during training that involved simply pressing all the available shapes until inevitably finding the target shape. This more ‘automated’ strategy would account for the higher error rate and could potentially lead to the overall more time spent searching for the target shape. Furthermore, this interpretation would not necessarily be exclusive of learning.

General Discussion

In our Fade-to-Reveal paradigm we observed an interesting effect of training: subjects showed differential improvement in performance for the two shapes. Subjects revealed significant reduction in reaction times for B-shapes while showing no significant changes in reaction time for A-shapes. While we believe this ultimately suggests that the subjects were coming to learn the sequence and rule behind the FOS transition, we are not able to conclusively say whether it is implicit learning of the sequence leading to a failure of the explicit order-discrimination test or simply a lack of power driving insignificant results at test.

In an effort to create more explicit learning of the sequence, we created our Search-and-Find paradigm expecting to see more robust learning and thus better performance on the order-recall test. This paradigm differed from the Fade-to-Reveal paradigm in requiring subjects to actively press on each of the faded shapes until they inevitably pressed on the target shape for that trial. We expected to see that this increase in involvement by the subject would ultimately lead to better learning. However, we failed to observe significance in training phase performance or in the results of the test for this paradigm either. However, we interpreted the quadratic trend for performance in B-shapes during training to be indicative of early learning with subsequent regression in performance due to the onset of fatigue.

The significant results found in the Fade-to-Reveal training gives us confidence that our subjects are indeed learning but that we are unable to measure learning with our current tests. If the subjects are in fact implicitly learning the sequence pairs in our Fade-to-Reveal training, we can consider testing for recognition of order by strategically introducing violations of the sequence (i.e. illegal transitions). We would then expect the subject to show an increase in the time needed to respond to the unexpected shape.

While we did not obtain the predicted results our findings suggest that adult subjects were indeed able to learn the temporal regularities between A-shapes and B-shapes and the information we have obtained will lead us to explore a different testing paradigm. An indirect test for order recognition which utilizes the introduction of sequence violations, unbeknownst to the subject, could potentially be a better indicator of a subject's ability to recognize temporal regularities. Considering that this particular paradigm will eventually to be performed on younger children between the ages of 2 and 5 years of age, an implicit test based on reaction times is a much simpler test for young children. Another possibility is to record looking times to legal and illegal temporal order reminiscent of the types of tests conducted by Kirkham et al., 2002 and Bulf et al., 2011 which investigated visual statistical learning in newborns and infants. While there is more work remaining, we have made clear steps towards developing a paradigm that will reveal information about implicit learning of visual temporal order.

References

Bulf H., Johnson S. P., & Valenza E. (2011). Visual statistical learning in the newborn infant. *Cognition* 121, 127–13210.1016

Erickson CA., & Desimone R. (1999) Responses of macaque perirhinal neurons during and after visual stimulus association learning. *J. Neurosci*.19:10404–10416.

Hafting, T., Fyhn, M., Molden, S., Moser, M-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex". *Nature* 436(7052): 801–806

Jabès A, Banta-Lavenex P, Amaral DG, Lavenex P. (2011) Postnatal development of the hippocampal formation: A stereological study in macaque monkeys. *J Comp Neurol*. 519:1051–1070.

Johnson A. & Redish A.D. (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* 27:12176–12189

Kirkham, N.Z., Slemmer, J.A., & Johnson, S. P. (2002). Visual statistical learning in infancy: evidence of a domain general learning mechanism. *Cognition, 83(2)*, B35-B42.

Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature*, 335:817-820.

O'Keefe J, & Dostrovsky. J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* 34(1): 171–175.

Ribordy F, Jabes A, Banta-Lavenex P., & Lavenex P. (2012) Development of allocentric spatial memory abilities in children from 18 months to 5 years of age. *Cognitive Psychology* 66:1–29.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.

Schapiro AC, Kustner LV, & Turk-Browne NB. (2012) Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*. 22:1622–1627.

Teinonen, T., Fellman, V., Näätänen, R., Alku, P., & Huotilainen, M. (2009). Statistical language learning in neonates revealed by event-related brain potentials. *BMC Neuroscience*, *10*, 21. doi:10.1186/1471-2202-10-21