Visual Statistical Learning in Infants

Natasha Zoe Kirkham (nzk2@cornell.edu)
Department of Psychology, Cornell University
Ithaca, NY 14853 USA

Jonathan Andrew Slemmer (jas234@cornell.edu)
Department of Psychology, Cornell University
Ithaca, NY 14853 USA

Scott P. Johnson (sj75@cornell.edu)
Department of Psychology, Cornell University
Ithaca, NY 14853 USA

Abstract

Statistical probability theory posits that we learn about regularly-occurring events in the perceptual environment by determining the likelihood of each event's occurrence (Aslin, Saffran, & Newport, 1998). The current study investigates infants' ability to extract properties of repetitive visual events and represent predictable combinations of visual elements. Using a novelty-preference paradigm, 2-, 5-, and 8-month-old infants were habituated to a continuous stream of colored shapes that were presented in a statistically predictable pattern, and then tested alternatively on the same sequence and a randomly-ordered sequence. The randomly-ordered sequence differed from the originally presented sequence only in between-shape transitional probabilities. At each age, infants demonstrated a significant novelty preference for the random sequence. In conjunction with Marcus, Vijayan, Rao, and Vishton (1999) and Saffran, Aslin, and Newport's (1996) work looking at statistical learning in language with 7- and 8-month-olds, these results can be taken as preliminary evidence of a domain general learning mechanism.

Introduction

One of the fundamental questions asked by developmental psychologists concerns how infants learn so much in so little time, with apparently very little explicit instruction. Research suggests that, as adults, we are remarkably good at implicit learning (e.g., see Stadler & Frensch, 1998 for a review of the implicit learning literature). Implicit learning can be defined as non-conscious facilitation of task performance due to information acquired during previous exposure. Given the robust nature of implicit learning skills in adults, it is perhaps a reasonable assumption that these skills play some role in early learning. The research on implicit learning in children suggests that they too show implicit sequence learning to the same degree as adults (Meulemans, Van der Linden, & Perruchet, 1998; Thomas, 1998). In other words, children showed increased reaction time in a task that contained a predictable sequence, and, like adults, did not have explicit knowledge of this sequence. There were no reports in the literature, however, with participants younger than 3 years of age until recent studies of statistical learning, a form of implicit learning based on statistical regularities in the perceptual environment.

Saffran, Aslin, and Newport (1996) and Alsin, Saffran, and Newport (1998) presented evidence that 8-month-old infants determine the statistical probability of neighboring speech sounds based on a 2-minute exposure. Infants heard four three-syllable "words" composed of 12 unique syllables (e.g., *tupiro*, *golabu*, *dapiku*, and *tilado*), presented in a continuous stream in random order (e.g., *dapikutupirotiladogolabutupiro* ...). Between-word spaces were removed, as were all other cues to word boundaries (e.g., rhythm, intonation, and stress). Thus, the only cues to word boundaries were the transitional probabilities between syllable pairs. For example, the transitional probability of *tu-pi* in this corpus is 1.00, because *pi* always follows *tu* within the word *tupiro*, whereas the probability of *ro-go* is .33, because *golabu* is one of three words that can follow *tupiro*.

After exposure, Saffran et al. (1996) presented infants with both familiar words from the corpus and "nonwords." Nonwords were created by combining the last syllable of one word with the first two syllables of a second word (e.g., *rogola* and *butupi*). Infants showed greater interest in the nonwords than in the words. On the logic that infants often exhibit a post-familiarization novelty preference (Bornstein, 1985), these results suggest that they detected the difference between words and nonwords. This outcome is necessarily based on learning of the transitional probabilities defining the stimuli.

This finding is consistent with a powerful statistical learning mechanism that supports language acquisition, and gives rise to questions concerning the generality of such mechanisms. It is possible, for example, that other kinds of knowledge are gained during infancy by learning statistical regularities in the environment. Indeed, Saffran, Johnson, Aslin, and Newport (1999) presented evidence that both adults and 8-month-old infants can perform the same statistics when presented with "words" consisting of non-linguistic tone sequences. These data suggest that statistical learning is not just a linguistic mechanism, allowing us to parse words from noise. If this is the case, then perhaps statistical learning is a mechanism that bolsters all sorts of learning in many different domains. Perhaps statistical learning is a domain general mechanism.

Following from this hypothesis the next logical step is to look at statistical learning in a non-auditory domain. The visual domain provides a lot of opportunities for patterns to present themselves, and the current study tested infants' ability to pick out statistical regularities in a stream of visual events. The data present evidence that 2-, 5-, and 8-month-old infants are able to extract properties of repetitive visual events in a way that allows them to represent a predictable combination of visual elements. Results from this series of studies showed that at each age, infants demonstrated a significant novelty preference for a visual sequence that differed from the originally presented sequence only in between-stimulus transitional probabilities.

Method

Participants

A total of 48 infants participated in this study: 16 infants at 2 months of age (M = 1.95 months), 16 infants at 5 months of age (M = 5.10 months) and 16 infants at 8 months of age (M = 7.99 months). Infants were recruited through a database of infants in the Ithaca, NY area. Informed consent was obtained from all parents, and the infants received a small toy or T-shirt as thanks for participation. All infants were full term and healthy. In addition to the 48 infants included in data analyses an additional 16 infants were tested but were not useable. Eight of these infants were 2-month-olds who either fell asleep during testing or were so fussy that looking times could not be judged correctly. In the 5- and 8-month-old group, data from one infant was not included due to equipment failure; the remaining seven infants were not included because they were so fussy that testing had to be terminated prior to the presentation of test trials.

Stimuli and Apparatus

Stimuli consisted of six looming colored-shapes (pink diamond, red octagon, yellow circle, blue cross, turquoise square, and green triangle) presented on a 53 cm computer monitor. The six colored-shapes were vector shapes that loomed from 2.35˚ to 14.59˚ of visual angle. Each stimulus loomed from 4 cm (2.35˚ visual angle) to 24 cm (14.59˚) in 1000 ms. There was no pause between stimulus presentations. The stimuli appeared in a continuous stream of randomly-ordered pairs (e.g., Pair 1: turquoise square followed by blue cross; Pair 2: yellow circle followed by pink diamond, Pair 3: green triangle followed by red octagon), with only transitional probabilities defining between-stimulus boundaries (see Figure 1 for an example of one shape sequence). The transitional probability within pairs was 1.0 and between pairs was 0.33 (see Figure 2 for an example of the transitions). In other words, for an individual infant, the pairs were always the same, but the order of the pairs within the sequence was random (e.g., turquoise square, blue cross, yellow circle, pink diamond, green triangle, red octagon, yellow circle, pink diamond, green triangle, red octagon, turquoise square, blue cross, turquoise square, blue cross….)

Procedure

Infants were tested individually and sat on a parent's lap 93 cm from the television monitor. The parent was instructed not to pay attention to their baby or to watch the screen. Infants were first shown the original stimulus sequence until looking declined according to a preset habituation criterion, determined by a sliding window algorithm that calculated over the course of a block of four trials when the infant's looking time had decreased by 50% from baseline. After habituation, the infants viewed six test displays alternating between familiar sequences, composed of the same three colored shape pairs, and novel sequences, produced by random recombinations of the same colored-shapes. The only difference, therefore, between the familiar and the novel sequences was that in the familiar sequences the first member of a colored shape pair predicted the second member, whereas in the novel sequences the colored shapes had no predictive value. There were six test trials in total (three familiar, and three novel). Test trials were counterbalanced across infants so that half the infants saw a familiar trial first and half the infants saw a novel trial first. The actual structure of the pairs was randomized across infants, so that it was very unlikely that any two infants saw the exact same pair sequence.

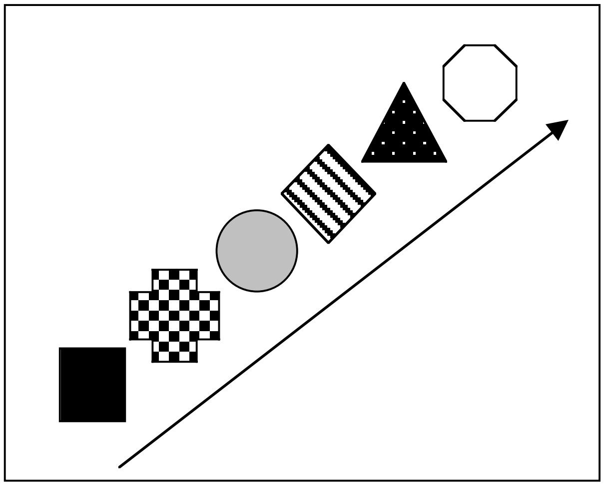

Figure 1: Example of a Shape Sequence (the actual shapes had unique colors, not black and white patterns).

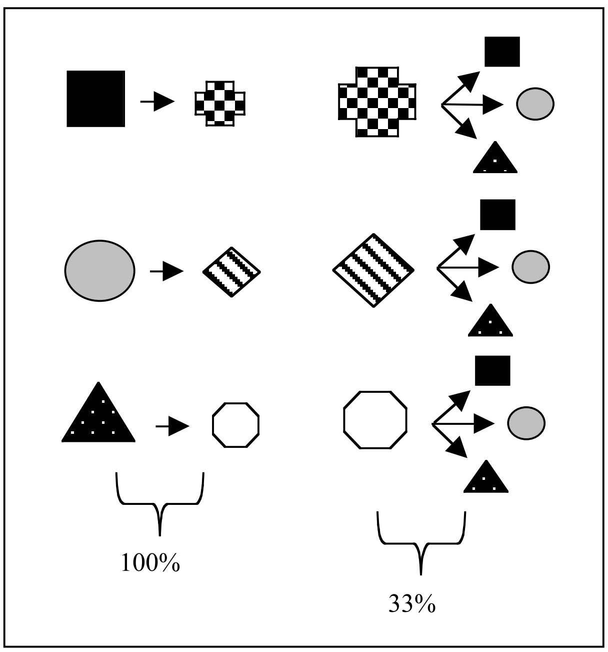

Figure 2: Example of the transitional probabilities between shapes.

Results

We hypothesized that infants would look longer at the novel sequences, if they were able to extract the visual statistical regularities available in the habituation stimuli.

The dependent measure was looking time at the familiar vs. the novel sequences. Looking time data in some cells were positively skewed (i.e., there were a few extremely long looking times); therefore, all data were log-transformed prior to analysis. A 3 (age group: 2-, 5-, or 8-month-olds) x 2 (test display: habituated vs. random sequence) ANOVA yielded a significant main effect of age, $F(2, 45) = 26.19$, $p <$ 001, the result of longer looking overall by the youngest infants (see Table 1). As was predicted, there was a reliable main effect of test display, $F(1, 45) = 17.89$, $p < .001$, the result of longer looking at the random test display (see Table 1). Planned comparisons showed that at each age infants looked longer at the random test display than at the familiar test display (at 2 months of age, $F(1, 45) = 5.47$, $p = .024$; at 5 months of age, , $F(1, 45) = 6.68$, $p = .013$; at 8 months of age, $F(1, 45) = 5.77$, $p = .02$; see Figure 3).

Table 1: Mean looking time in sec per test display according to age group.

Age Group	Familiar Sequence	Novel Sequence
2 months	23.67	32.95
5 months	8.10	11.39
8 months	6.07	9.28

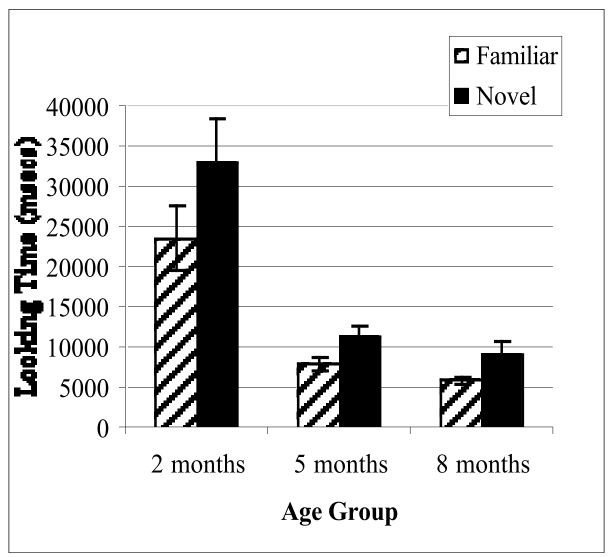

Figure 3: Infants' Looking Time to Familiar and Novel Patterns

Discussion

These results reveal that even very young infants are capable of picking up on statistical regularities in the visual environment. Moreover, they are discriminating between two visual sequences that differ *only* in these statistical probabilities. The sequences offered no cues as to the pairings: there were no pauses between stimulus presentations, the presentations were the same length, and the colored-shapes were not salient objects. Yet, each age group showed significantly longer looking time to the novel sequences. This evidence that statistical learning is evident in young infants, in the visual domain, lends support to the hypothesis that this could be an early developing mechanism that supports learning across domains.

It is, of course, worth noting that this is only the first step in addressing the ways in which learning develops. We have yet to determine exactly how infants are encoding the information. Is it pure statistics? Are babies taking note of co-occurrences of events and stimuli, and judging the probabilities of those co-occurrences or are they abstracting these statistics into higher-order rules? Research on language learning has suggested that developmental constraints on learning actually provide the necessary prerequisites for later complex learning (Elman, 1993; Newport, 1988, 1990). Our research offers the same suggestion; early learning in the visual domain starts with the acquisition of the simplest patterns, and over development more complex patterns are acquired. This domain general learning could provide the beginnings of domain specific learning: as learning becomes more complex, the types of learning differentiate according to domain

As adults, we are brilliant implicit learners, and this learning has very little decay and requires very little effort. As infants, all we do is look, listen and touch; we absorb information from the environment and quickly pick up on the causal relationships we experience around us. For example, it does not take an infant very long to figure out that crying creates a desired effect, that of a parent's presence, and that the presence of a parent tends to predict food or comfort. This type of associative learning seems natural and adaptive. But, what if all initial learning is associative, and not dependent on a conditioned reward/punishment outcome. Note that the point here is to focus on initial learning, not to suggest that all learning develops in the same way. Perhaps the associations are all that is needed to elicit a type of implicit learning. Admittedly, the associations that are relevant and salient and do produce pleasant or unpleasant end results might be learned faster and remembered longer. We do not think that the infants we tested are going to predict the arrival of a turquoise square every time they see a yellow circle, for example, but it is interesting to observe that the associations were there at least long enough for the babies to notice a difference when the turquoise square did not predict a yellow circle. What is most interesting about the statistical learning process is not that it may be domain general but that it seems to work in situations that do not have reward/punishment end results, and therefore, seems capable of supporting a great deal of initial knowledge acquisition.

Acknowledgments

The authors would like to thank the many parents and infants for their participation, and all the members of the Cornell Baby Lab for help with infant recruitment. As well, we would like to thank Richard Aslin for many helpful comments. The research was supported by NSF grant BCS-9910779.

References

Aslin, R. N.; Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological-Science, 9,* 321-324.

Bornstein, M. H. (1985). Habituation of attention as a measure of visual information processing in human infants: Summary, systematization, and synthesis. In G. Gottlieb & N. A. Krasnegor (Eds.), *Measurement of audition and vision in the first year of postnatal life: A methodological overview*. Norwood, NJ: Ablex.

Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition, 48,* 71-99.

Marcus, G.F., Vijayan, S, Rao, S.B., & Vishton, P.M. (1999). Rule learning by seven-month-old infants. *Science*, *283(*5398), 77-80.

Meulemans, T, Van der Linden, M., & Perrouchet, P. (1998). Implicit learning in children. *Journal of Experimental Child Psychology, 69(*3), 199-221.

Newport, E. L. (1988). Constraints on learning and their role in language acquisition: Studies of the acquisition of American Sign Language. *Language Sciences, 10,* 147-172.

Newport, E. L. (1990). Maturational constraints on language learning. *Cognitive Science, 14,* 11-28.

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

Saffran, J.R, Johnson, E. K., Aslin, R.N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, *70(*1), 27-52.

Spelke, E. S., & Van de Walle, G. (1993). Perceiving and reasoning about objects: Insights from infants. In N. Eilan, R. A. McCarthy, & B. Brewer (Eds.), *Spatial representation: Problems in philosophy and psychology*. Oxford: Blackwell.

Stadler, M.A., & Frensch, P.A. (Eds.) (1998). *Handbook of implicit learning.* Thousand Oaks: Sage Publications

Thomas, K.M. (1998). Behavioral and electrophysiologic measures of implicit learning in children. *Unpublished doctoral dissertation.*