Journal of Experimental Child Psychology 133 (2015) 47–56

Contents lists available at ScienceDirect

Journal of Experimental Child Psychology

journal homepage: www.elsevier.com/locate/jecp

Infants' statistical learning: 2- and 5-month-olds' segmentation of continuous visual sequences

Lauren Krogh Slone *, Scott P. Johnson

Department of Psychology, University of California, Los Angeles, Los Angeles, CA 90095, USA

A R T I C L E I N F O

Article history:
Received 12 May 2014
Revised 14 January 2015

Keywords:
Visual statistical learning
Sequence segmentation
Transitional probability
Co-occurrence frequency
Cognitive development
Habituation

A B S T R A C T

Past research suggests that infants have powerful statistical learning abilities; however, studies of infants' visual statistical learning offer differing accounts of the developmental trajectory of and constraints on this learning. To elucidate this issue, the current study tested the hypothesis that young infants' segmentation of visual sequences depends on redundant statistical cues to segmentation. A sample of 20 2-month-olds and 20 5-month-olds observed a continuous sequence of looming shapes in which unit boundaries were defined by both transitional probability and co-occurrence frequency. Following habituation, only 5-month-olds showed evidence of statistically segmenting the sequence, looking longer to a statistically improbable shape pair than to a probable pair. These results reaffirm the power of statistical learning in infants as young as 5 months but also suggest considerable development of statistical segmentation ability between 2 and 5 months of age. Moreover, the results do not support the idea that infants' ability to segment visual sequences based on transitional probabilities and/or co-occurrence frequencies is functional at the onset of visual experience, as has been suggested previously. Rather, this type of statistical segmentation appears to be constrained by the developmental state of the learner. Factors contributing to the development of statistical segmentation ability during early infancy, including memory and attention, are discussed.

* Corresponding author.
E-mail address: laurenkslone@gmail.com (L.K. Slone).

http://dx.doi.org/10.1016/j.jecp.2015.01.007

Introduction

A central question of cognitive science and developmental psychology concerns the mechanisms underlying learners' ability to transform visual and auditory stimuli from the environment into mental representations. This is a particularly challenging task during infancy, when so much of the environment is unfamiliar. One important learning mechanism available to infants is the ability to segment continuous input based on statistical relations among elements in that input. Because this "statistical learning" ability has been posited to underlie cognitive skills ranging from infants' language acquisition to apprehension of event structure (e.g., Bates & Elman, 1996; Roseberry, Richie, Hirsh-Pasek, Golinkoff, & Shipley, 2011), elucidating the mechanisms of and constraints on statistical learning ability during infancy is critical.

Most research investigating infants' statistical segmentation ability has employed auditory input (e.g., Aslin, Saffran, & Newport, 1998; Pelucchi, Hay, & Saffran, 2009; Saffran, Aslin, & Newport, 1996; Saffran, Johnson, Aslin, & Newport, 1999); relatively little is known about infants' segmentation of visual sequences (to our knowledge, only six published articles: Bulf, Johnson, & Valenza, 2011; Kirkham, Slemmer, & Johnson, 2002; Kirkham, Slemmer, Richardson, & Johnson, 2007; Marcovitch & Lewkowicz, 2009; Roseberry et al., 2011; and Stahl, Romberg, Roseberry, Golinkoff, & Hirsh-Pasek, 2014). Interestingly, these studies offer differing accounts of the developmental trajectory of infants' visual statistical learning; the findings of Bulf et al. (2011) and Kirkham et al. (2002) suggest that visual statistical learning ability is in place by birth and shows no signs of developmental change between 2 and 5 months of age, but the results of Marcovitch and Lewkowicz (2009) suggest that visual statistical learning undergoes considerable development between these ages.

These two sets of studies, however, differed in how statistical cues signaled sequence structure. Whereas Kirkham et al.'s (2002) sequences contained *redundant* co-occurrence frequency and transitional probability (TP) information, Marcovitch and Lewkowicz (2009) tested for the *independent* contributions of co-occurrence frequency and TP information. Co-occurrence frequency is defined in these studies as the rate at which two items (X and Y) appear consecutively in a sequence, whereas TP is the conditional probability of item Y following item X, given that X has appeared, and thus is a measure of the strength with which X predicts Y. One possible explanation of 2-month-olds' differential success in the Kirkham et al. (2002) and Marcovitch and Lewkowicz (2009) experiments, therefore, is that the redundant statistical cues used by Kirkham and colleagues facilitated 2-month-olds' visual statistical learning. The aim of the current research was to test this prediction. Specifically, the current study examined whether, with redundant TP and co-occurrence frequency cues, 2-month-olds would show evidence of statistical segmentation of continuous visual sequences.

Visual statistical learning: Structured versus random sequences

Kirkham et al. (2002) were the first to examine infants' statistical learning of visual sequences. Infants (2-, 5-, and 8-month-olds) viewed a continuous stream of six colored shapes that appeared one at a time on a computer monitor. Shapes were organized into three pairs. Pairs were not allowed to repeat, such that pair boundaries were marked by differing TPs between shapes within pairs (1.0 TP) and between pairs (.50 TP). After habituating to this sequence, infants at all three ages looked longer at a randomly ordered sequence composed of the same six shapes relative to the familiar sequence, suggesting that discrimination between TP-structured and nonstructured visual sequences is accomplished by infants as young as 2 months. Bulf et al. (2011) used a similar method to investigate visual statistical learning in newborn infants. Infants were shown either four or six shapes organized into pairs. In contrast to Kirkham et al. (2002), pairs were allowed to repeat in Bulf et al. (2011), such that TPs between pairs were .33 in the six-shape condition and .50 in the four-shape condition. Bulf and colleagues found that newborns could discriminate between the familiar habituation sequence and a randomly ordered sequence with the set of four shapes (but not the set of six shapes). These studies suggested that infants are adept at detecting statistical regularities in input from the visual modality and, together with research documenting infants' statistical segmentation of auditory sequences, were taken as evidence of the domain generality of statistical segmentation ability.

However, although Kirkham et al. (2002) and Bulf et al. (2011) demonstrated that infants between birth and 8 months of age discriminate random sequences from structured ones, it is unclear what information infants used to make this discrimination. A recent study by Addyman and Mareschal (2013) found spontaneous preferences by 5-month-olds for random versus structured sequences of looming shapes in which pairs were allowed to repeat. The authors suggested that infants were responding to local redundancies engendered by repetitions of shape pairs in the structured sequence rather than responding to the global statistical structure. This suggests that infants' preference for the random test sequence in Bulf et al. (2011) might not have reflected segmentation based on TP statistics. Moreover, TP and co-occurrence frequency perfectly covaried in Kirkham et al. (2002) and Bulf et al.(2011). If infants were tracking the overall statistical structure of the habituation sequence, they could have used either or both of these statistics to distinguish between the test sequences. Identifying whether infants are in fact capable of segmenting continuous visual sequences, and what statistics infants may use to accomplish this, is fundamental to identifying both the mechanisms of and constraints on visual statistical learning as well as to assessing domain generality.

Visual statistical learning: Segmentation

Marcovitch and Lewkowicz's (2009) study with 2.5-, 4.5-, and 8.5-month-olds came closer to achieving this goal. Similar to Kirkham et al. (2002), Marcovitch and Lewkowicz (2009) habituated infants to a continuous stream of six colored shapes organized into pairs. In contrast to Kirkham et al. (2002), however, Marcovitch and Lewkowicz's (2009) design manipulated co-occurrence frequency information independent of TP and assessed sensitivity to both features. Infants were presented with two types of pairs during habituation: pairs with high co-occurrence frequency but low TP and pairs with low co-occurrence frequency but high TP. Importantly, their design assessed segmentation based on global statistical information given that test trials compared infants' looking to repeated presentations of statistically probable and improbable pairs of items, not continuous sequences. Detection of TPs was examined by comparing looking to a high-TP test pair with the average looking to two other low-TP test pairs. Similarly, detection of co-occurrence frequency information was examined by comparing looking to a high-frequency test pair with the average looking to two low-frequency test pairs. Interestingly, Marcovitch and Lewkowicz found evidence for recognition of both TP-defined and frequency-defined units by 4.5- and 8.5-month-olds but not by 2.5-month-olds. This contrasts with Kirkham et al.'s (2002) finding that 2-month-olds successfully distinguish random sequences from structured ones and raises the question of why 2.5-month-olds provided no evidence of statistical learning in Marcovitch and Lewkowicz's (2009) study.

Marcovitch and Lewkowicz (2009) ascribed the result to the greater difficulty of their task given that it assessed segmentation based on differences in TP and co-occurrence frequency rather than discrimination of patterned sequences from random ones. However, we reasoned that 2.5-month-olds' failure to show evidence of segmentation in Marcovitch and Lewkowicz's study may have resulted from a lack of redundant statistical cues (i.e., co-occurrence frequency and TP information) to support learning rather than from an inability to segment continuous visual sequences per se. That is, redundant correlated cues are particularly helpful to children's learning (e.g., Lewkowicz, 2004; Thiessen, 2012; Yoshida & Smith, 2005). Moreover, infants have been found to succeed in rule learning and statistical learning tasks at an earlier age when provided with redundant cues to sequence structure (e.g., Frank, Slemmer, Marcus, & Johnson, 2009; Kirkham et al., 2007) compared with when redundant cues are lacking. One explanation of this phenomenon is that, when presented with a sequence of input, infants generate a set of hypotheses about the underlying structure organizing the sequence (e.g., two pairs, three triplets). The input is compared with these hypotheses, and hypotheses consistent with the input are credited as more likely (Gerken, 2006; Tenenbaum & Griffiths, 2001). Redundant cues may facilitate correct segmentation because the likelihood of multiple cues converging on a single segmentation hypothesis by chance is lower than the likelihood of a single cue doing so (Frank et al., 2009). We reasoned, therefore, that redundant co-occurrence frequency and TP information together might facilitate younger infants' performance in the current task and would imply that the mechanisms responsible for visual segmentation are present earlier than suggested by Marcovitch and Lewkowicz (2009).

The goal of the current study was to examine whether 2-month-olds segment continuous visual sequences when provided with redundant statistical cues to unit boundaries. We exposed 2- and 5-month-olds to a continuous stream of colored shapes in which both TP and co-occurrence frequency were higher between shapes within units relative to across unit boundaries. If the ability to segment visual sequences based on these statistics does not emerge until later than 2 months of age, only 5-month-olds should discriminate between units (high TP, high co-occurrence frequency) and part-units (low TP, low co-occurrence frequency). However, if the ability to segment visual sequences based on TP and co-occurrence frequency information is intact at an earlier age but requires these cues to converge on the same unit boundaries, both 2- and 5-moth-olds should discriminate between units and part-units.

Method

Participants

A sample of 20 healthy full-term 2-month-olds (M_{age} = 2 months 12 days, range = 1;27 to 2;20 [months;days], 11 female and 9 male) and 20 healthy full-term 5-month-olds (M_{age} = 5 months 1 day, range = 4;15 to 5;15, 6 female and 14 male) were tested using an infant-controlled habituation technique. Data from an additional 21 infants were excluded from the final sample due to fussiness or disinterest (5 2-month-olds and 4 5-month-olds), falling asleep (1 2-month-old), failure to habituate (3 2-month-olds and 4 5-month-olds), parental interference (1 5-month-old), or failure to observe the entire two-shape test sequences (1 2-month-old and 2 5-month-olds). All infants included in the final analyses were alert and engaged during both learning and test. Infants were recruited by letter and telephone from hospital records and were given a small gift (a toy or baby T-shirt) for their participation.

Apparatus and stimuli

A Dell OptiPlex 755 computer and 61.5-cm color monitor were used to present stimuli and collect looking time data. Stimuli consisted of six colored shapes (red circle, orange diamond, white heart, green plus, blue square, and yellow triangle) presented one at a time for 750 ms and looming from 1.5 to 5.5 cm in height ($\sim$1.43°–5.25°) in front of a black background. Each habituation and test sequence was presented until the infant visually fixated away from the monitor for more than 2 s or until 90 s of looking had accumulated, as coded by the primary experimenter. A second judge who was blind to the hypothesis and which trial was presented recoded the looking times of 35% of the infants from digital recordings of the sessions. The Pearson product–moment pairwise correlation between the two coders was greater than .94, and the mean difference between the two coders' scores was less than 0.30 s ($p > .60$). Prior to the first habituation trial and prior to every subsequent trial, an attention-getter was shown to attract the infant's gaze back to the screen, at which time the next trial was started immediately.

Procedure

Each infant sat on a caretaker's lap approximately 60 cm from the computer monitor. The habituation sequence consisted of three psuedo-randomly ordered shape pairs (e.g., Pair 1: orange diamond and blue square; Pair 2: yellow triangle and red circle; Pair 3: white heart and green plus) (Fig. 1A). Pairs could not repeat, and there were no breaks or delays between shapes or pairs, such that the only cues to shape pairings were interstimulus TPs (1.0 within pairs, .50 between pairs) and co-occurrence frequencies (each pair: 16.67%; each part-unit: 8.33%).

Following habituation, infants were presented with three repetitions of two unique test trials in alternation for a total of six test trials. Each test trial consisted of a repetition of two shapes separated by a 750-ms pause to mark unit boundaries. The familiar sequence was a pair from habituation (TP = 1.0, co-occurrence frequency = 16.67%) (Fig. 1B), whereas the 2–1 sequence spanned a pair

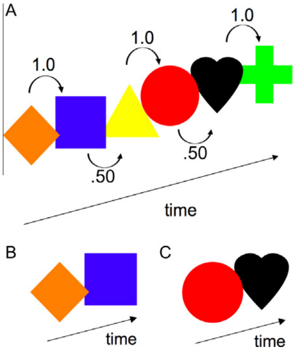

Fig. 1. Sample habituation sequence (A), familiar test sequence (B), and 2–1 test sequence (C). Numbers above arrows indicate TPs between shapes. The heart is shown in black here but was white in the experiment.

boundary (the second member of a habituation pair, followed by the first member of a different habituation pair; TP = .50, co-occurrence frequency = 8.33%) (Fig. 1C). Which test trial infants viewed first was randomized. The only differences between the two test displays were the TPs between and co-occurrence frequencies of the shapes; this ensured that any looking time difference in the two types of test trials would necessarily be related to the statistical nature of the sequences. If infants segment the habituation sequence on the basis of TP and/or co-occurrence frequency information, they should discriminate between the familiar and 2–1 test sequences.

Results and discussion

The goal of the current study was to examine whether, with redundant TP and co-occurrence frequency cues to unit boundaries, 2-month-olds are capable of segmenting continuous visual sequences. The experiment directly tested segmentation by requiring infants to discriminate a statistically probable pair from an improbable pair at test. If the ability to segment visual shape sequences based on TP and co-occurrence frequency information does not emerge until later than 2 months of age, only the 5-month-olds should discriminate between the familiar and 2–1 test sequences. However, if the ability to segment these types of visual sequences is intact at 2 months but simply requires converging TP and co-occurrence frequency cues to unit boundaries at this younger age, we hypothesized that both 2- and 5-month-olds should discriminate between the familiar (TP = 1.0, co-occurrence frequency = 16.67%) and 2–1 (TP = .50, co-occurrence frequency = 8.33%) test sequences.

To examine infants' looking behavior during test, the six test trials were divided into three blocks, each consisting of two types of sequences: a familiar sequence and a 2–1 sequence. A $3 \times 2 \times 2$ mixed analysis of variance (ANOVA), with block and sequence type as within-participants factors and age as a between-participants factor, revealed a significant main effect of block, $F(2, 38) = 5.59$, $p < .01$, partial $\eta^2 = .13$. As can be seen in Fig. 2, infants' looking times decreased across the three blocks of test trials, indicating that infants became familiar with the overall experimental setup. There was also a main effect of age, $F(1, 38) = 9.52$, $p < .01$, partial $\eta^2 = .20$, the result of longer looking overall by the 2-month-olds. (Very young infants often exhibit longer looking times in visual tasks compared with older infants, which may reflect developmental differences in infants' information processing skill [see Johnson, 1996].) Critically, the ANOVA also revealed a significant sequence type by age interaction, $F(1, 38) = 4.41$, $p < .05$, partial $\eta^2 = .10$. The 5-month-olds looked significantly longer to the 2–1 test sequence compared with the familiar test sequence, $t(19) = 2.76$, $p < .05$, $d = 0.62$, two-tailed, whereas the 2-month-olds did not, $t(19) = 0.87$, $p = .40$, $d = 0.20$, two-tailed. There was no interaction with block, $F(2, 38) = 0.97$, $p = .38$, partial $\eta^2 = .02$, indicating that infants' preference (or lack of preference) was consistent across the three blocks.

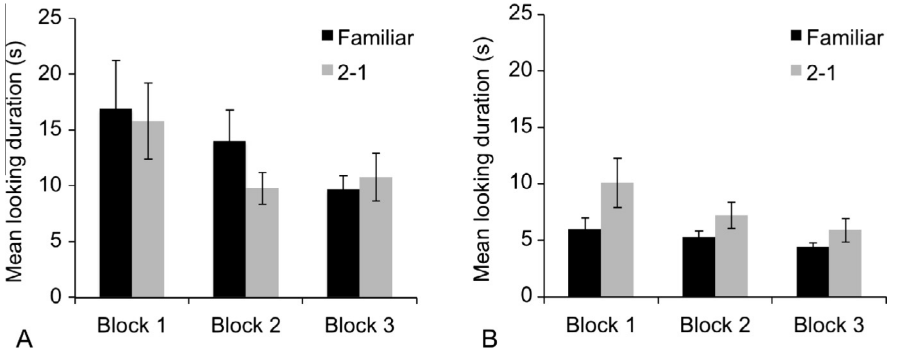

Fig. 2. The 2-month-olds' (A) and 5-month-olds' (B) mean looking durations to the familiar and 2–1 test displays by block. Error bars represent standard errors.

These findings suggest considerable development of statistical segmentation ability between 2 and 5 months of age. Because TP and co-occurrence frequency provided redundant sources of information for segmentation in the current study, we cannot determine whether 5-month-olds' successful performance was based on sensitivity to one or both of these statistics. However, using a frequency balanced design, Fiser and Aslin (2002) demonstrated that infants discriminate high-TP shape pairs from low-TP ones that occur equally frequently by at least 9 months of age. Moreover, Marcovitch and Lewkowicz's (2009) findings suggest that 5-month-olds respond to TP and co-occurrence frequency statistics independent of one another. These findings suggest that 5-month-olds in the current study may have been sensitive to differences in both TP and co-occurrence frequency and may have used either statistic successfully to segment the sequence.

The 2-month-olds we observed, in contrast, provided no evidence of sensitivity to either TP or co-occurrence frequency information in the visual sequence. This is in line with Marcovitch and Lewkowicz's (2009) results with 2-month-olds but contrasts with Kirkham et al.'s (2002) finding that 2-month-olds can discriminate TP-structured sequences from random ones. Thus, contrary to our hypothesis, the redundant TP and co-occurrence frequency information inherent in both Kirkham et al. (2002) and the current study was insufficient here to facilitate learning in 2-month-olds.

One possible explanation for these results is that, similar to Marcovitch and Lewkowicz (2009), the current segmentation task was overly complex and taxing for these young infants, inhibiting performance. However, 2-month-olds' low level of attrition due to fussiness or disinterest (16% vs. 52% in Marcovitch & Lewkowicz's experiment) suggests that this is not the primary reason for the apparent inability of 2-month-olds to segment visual sequences under tested conditions.

We believe a more probable explanation is that infants at this young age lack the cognitive processing abilities necessary for statistical segmentation based on TP and co-occurrence frequency information given these types of materials (see Bulf et al., 2011). The current experiment did not directly examine the mechanisms underlying visual statistical segmentation, yet it is informative to consider the differences between how learning was tested in the current study and in Kirkham et al.'s (2002) study. Kirkham and colleagues presented infants with structured and random test sequences that included all six shapes from habituation with no pauses between items. In contrast, the current experiment tested segmentation more strictly by presenting repetitions of one shape pair and one part-pair from habituation, with boundaries demarcated by 750-ms pauses. Consequently, the current experiment required three related perceptual and cognitive skills that Kirkham et al.'s study did not: (a) computation of TPs and/or co-occurrence frequencies during habituation, which inherently required (b) attention to more shapes, and (c) recollection of that information during test trials.

For instance, infants could have discriminated between Kirkham et al.'s (2002) test sequences on the basis of statistical relations (e.g., co-occurrence frequency, TP) between pairs of shapes, but they could also have discriminated the sequences by attending to frequencies of individual shapes, a much less demanding task. That is, although shape frequency was balanced in Kirkham et a.'s structured test sequence (to produce specific TPs between shape pairs), the only constraint on the random test sequence was that individual shapes could not repeat. With test trials lasting less than 30 s on average (<30 shapes per trial), it is possible, if not likely, that the frequencies of individual shapes differed across the two types of test trials. Thus, it is possible that infants' preference for the random sequence in Kirkham et al.'s study was a result of sensitivity to differences in shape frequency between the two test sequences rather than a result of sensitivity to statistical relations between shapes during habituation. Although this would imply sensitivity to some types of statistical structures (individual shape frequencies), it would nevertheless suggest that at 2 months of age infants might not track the more complex TP and co-occurrence frequency statistics that infants track by 5 months of age. One way to test this hypothesis would be to examine whether 2-month-olds discriminate structured test sequences from random ones as in Kirkham et al.'s study when individual shape frequencies are equated across the two types of test sequences. Another interesting approach would be to examine whether exposure to Kirkham et al.'s random sequence during habituation or no habituation exposure at all (see Addyman & Mareschal, 2013) nevertheless leads to a preference for random test sequences compared with structured test sequences.

Although such investigations might clarify whether infants were in fact tracking statistical relations between shapes during habituation in Kirkham et al.'s (2002) study, it would leave unanswered the critical question of what exactly is underlying the development of visual statistical segmentation ability between 2 and 5 months of age. The 2-month-olds provided no evidence under tested conditions to discriminate between test sequences in both the current study and the Marcovitch and Lewkowicz (2009) study, and this may have resulted from a failure of segmentation, a failure of memory for the segmented sub-sequences, or both. Currently, the development of memory and attention in such young infants is a topic of limited research, making it difficult to identify specific memory or attentional processes potentially responsible for the shift in statistical segmentation performance between 2 and 5 months of age. Nevertheless, several possibilities have been suggested.

Bulf et al. (2011) suggested that increases in short-term memory capacity during early infancy may facilitate the development of statistical learning. Detection of statistical relations among sequential items necessitates memory for previous items. This is perhaps most applicable to visual statistical learning tasks in which stimulus onset intervals are typically much longer than those of auditory tasks, requiring visual items to be held in memory for greater durations. Visual short-term memory (VSTM) performance has been shown to increase roughly linearly across the first postnatal year (Diamond, 1985; see Bell & Morasch, 2007, for a review), suggesting that increases in VSTM may be an important factor underlying the development of statistical learning during infancy. However, this account leaves unexplained the finding of successful discrimination by 2-month-olds in Kirkham et al.'s (2002) study.

Another possibility relates to infants' allocation of attentional resources. Although statistical learning is largely implicit in nature (e.g., Meulemans, Van der Linden, & Perruchet, 1998; Saffran, Newport, Aslin, Tunick, & Barrueco, 1997), there is evidence supporting the role of attention in statistical learning (Emberson, Conway, & Christiansen, 2011; Toro, Sinnett, & Soto-Faraco, 2005; Turk-Browne, Junge, & Scholl, 2005). For instance, when adults observed two interleaved streams of shapes in two different colors but were instructed to attend to only one color, only the statistical relations in the attended shape stream were learned (Turk-Browne et al., 2005). Such findings suggest that simply viewing stimuli might not always be enough to allow encoding of statistical regularities. Thus, although both the 2- and 5-month-olds in the current study were alert and looking at the shape sequences, it is possible that differences in the 2- and 5-month-olds' allocation of attention may have played a role in the younger infants' inability to track TP and co-occurrence frequency statistics in the shape sequence.

One way to examine the role of attention in infants' statistical segmentation is to employ stimuli that capture infants' attention to varying degrees. The intersensory redundancy hypothesis (IRH) argues that synchronous cues from multiple modalities are more salient to infants than synchronous unimodal cues (see Bahrick & Lickliter, 2012). This increased salience is thought to direct infants' attention to the multimodal event, resulting in enhanced perception and discrimination (e.g., Bahrick & Lickliter, 2000; Bahrick, Lickliter, & Flom, 2004). Future research could examine whether redundant multisensory cues facilitate infants' statistical segmentation more so than redundant unimodal cues of the sort employed here.

In addition, research suggests that the attentional facilitation of multimodal redundancy is most pronounced in very young infants. For instance, in a study of serial order perception, Lewkowicz (2004) found that 4-month-olds perceived serial order only when events were specified redundantly in the visual and auditory modalities, whereas 8-month-olds were able to perceive serial order in events specified both unimodally and multimodally. Lewkowicz and Berent (2009) found that 4-month-olds learn serial order in such multimodal events by tracking statistical associations. Together, these findings suggest that multimodal cues to unit boundaries may facilitate statistical segmentation in very young infants. Future research should investigate this possibility to determine whether infants are in fact capable of statistical segmentation based on TP or co-occurrence frequency prior to 5 months of age given the right stimulus materials. One way to test this would be to add synchronous auditory information to the visual habituation sequence employed in the current experiment. The auditory cues could be syllables from an artificial language that are either uniquely paired with the shapes (Condition 1) or randomly paired with the shapes each time they occur (Condition 2). Test trials would assess infants' discrimination between the familiar and 2–1 test sequences presented only in the visual modality. Successful discrimination by 2-month-olds in Condition 1 would suggest that infants are in fact capable of statistical segmentation of visual sequences at this young age given redundant multimodal cues to unit boundaries. Successful discrimination by 2-month-olds in Condition 2 would suggest that the nature of the developmental difference observed in the current experiment may be attentional. That is, multimodal cues may recruit and direct infants' attention more so than the unimodal cues used in the current study and, thus, facilitate learning despite not providing additional evidence for segmentation.

Conclusions

The current study affirms the power of statistical learning in young infants. By 5 months of age, infants appear to be able to detect the TP and co-occurrence frequency structure of visual sequences after only a few minutes of exposure in a highly constrained unnatural setting. This finding, in conjunction with the findings of Fiser and Aslin (2002) and Saffran and colleagues (Pelucchi et al., 2009; Saffran et al., 1996, 1999), suggests that early development is highly attuned to TP and co-occurrence frequency statistics across a wide range of input. Nevertheless, the current study does not support the claim that segmentation based on TPs and/or co-occurrence frequencies is functional at the onset of visual experience (Bulf et al., 2011; Kirkham et al., 2002). The current study tested the hypothesis that redundant TP and co-occurrence frequency cues are what facilitated infants' learning in Kirkham et al. (2002). However, we did not find positive evidence of segmentation with 2-month-olds, suggesting that these redundant statistical cues are not enough to facilitate infants' statistical segmentation, at least given these types of materials. Rather, the ability to segment continuous visual sequences based on TP and/or co-occurrence frequency information appears to be constrained by the developmental state of the learner (Krogh, Vlach, & Johnson, 2013; Thiessen, 2010).

Although further research is needed to identify the specific constraints that may underlie the development of visual statistical segmentation ability during early infancy, the current findings have important implications for models of infant cognition in general and statistical segmentation in particular. Little is known about statistical segmentation during early infancy, yet a number of computational models of statistical segmentation appear to assume that it operates without attentional, memory, or computational constraints (e.g., Goldwater, Griffiths, & Johnson, 2009; Swingley, 2005). Such models will acquire greater plausibility with the inclusion of realistic capacity limitations (see Perruchet & Vinter, 1998 PARSER model for an example) and means of information acquisition (see Bulf et al., 2011, for a discussion).

Acknowledgments

Funding for this project was provided by the National Science Foundation Graduate Research Fellowship Program under Grant DGE-0707424 to L.K.S. and by NIH Grant R01-HD73535 to S.P.J. We thank Bryan Nguyen for his help with programming and thank the volunteers and staff of the

Baby Lab at the University of California, Los Angeles, for their help with participant recruitment. We also thank the parents and infants who participated in this study as well as three anonymous reviewers for excellent suggestions and critiques during the review process.

References

Addyman, C., & Mareschal, D. (2013). Local redundancy governs infants' spontaneous orienting to visual–temporal sequences. *Child Development, 84*, 1137–1144.
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.
Bahrick, L. E., & Lickliter, R. (2000). Intersensory redundancy guides attentional selectivity and perceptual learning in infancy. *Developmental Psychology, 36*, 190–201.
Bahrick, L. E., & Lickliter, R. (2012). The role of intersensory redundancy in early perceptual, cognitive, and social development. In A. Bremner, D. J. Lewkowicz, & C. Spence (Eds.), *Multisensory development* (pp. 183–205). Oxford, UK: Oxford University Press.
Bahrick, L. E., Lickliter, R., & Flom, R. (2004). Intersensory redundancy guides infants' selective attention, perceptual, and cognitive development. *Current Directions in Psychological Science, 13*, 99–102.
Bates, E., & Elman, J. (1996). Learning rediscovered. *Science, 274*, 1849–1850.
Bell, M. A., & Morasch, K. C. (2007). The development of working memory in the first 2 years of life. In L. M. Oakes & P. J. Bauer (Eds.), *Short and long-term memory in infancy and early childhood* (pp. 27–50). New York: Oxford University Press.
Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition, 121*, 127–132.
Diamond, A. (1985). Development of the ability to use recall to guide action, as indicated by infants' performance on AB. *Child Development, 56*, 868–883.
Emberson, L. L., Conway, C. M., & Christiansen, M. H. (2011). Timing is everything: Changes in presentation rate have opposite effects on auditory and visual implicit statistical learning. *Quarterly Journal of Experimental Psychology, 64*, 1021–1040.
Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National academy of Sciences of the United States of America, 99*, 15822–15826.
Frank, M. C., Slemmer, J. A., Marcus, G. F., & Johnson, S. P. (2009). Information from multiple modalities helps 5-month-olds learn abstract rules. *Developmental Science, 12*, 504–509.
Gerken, L. (2006). Decisions, decisions: Infant language learning when multiple generalizations are possible. *Cognition, 98*, B67–B74.
Goldwater, S., Griffiths, T. L., & Johnson, M. (2009). A Bayesian framework for word segmentation: Exploring the effects of context. *Cognition, 112*, 21–54.
Johnson, S. P. (1996). Habituation patterns and object perception in young infants. *Journal of Reproductive and Infant Psychology, 14*, 207–218.
Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition, 83*, B35–B42.
Kirkham, N. Z., Slemmer, J. A., Richardson, D. C., & Johnson, S. P. (2007). Location, location, location: Development of spatiotemporal sequence learning in infancy. *Child Development, 78*, 1559–1571.
Krogh, L., Vlach, H., & Johnson, S. P. (2013). Statistical learning across development: Flexible yet constrained. *Frontiers in Language Sciences, 3*, 598. http://dx.doi.org/10.3389/fpsyg.2012.00598.
Lewkowicz, D. J. (2004). Perception of serial order in infants. *Developmental Science, 7*, 175–184.
Lewkowicz, D. J., & Berent, I. (2009). Sequence learning in 4-month-old infants: Do infants represent ordinal information? *Child Development, 80*, 1811–1823.
Marcovitch, S., & Lewkowicz, D. L. (2009). Sequence learning in infancy: The independent contributions of conditional probability and pair frequency information. *Developmental Science, 12*, 1020–1025.
Meulemans, T., Van der Linden, M., & Perruchet, P. (1998). Implicit sequence learning in children. *Journal of Experimental Child Psychology, 69*, 199–221.
Pelucchi, B., Hay, J. F., & Saffran, J. R. (2009). Statistical learning in a natural language by 8-month-old infants. *Child Development, 80*, 674–685.
Perruchet, P., & Vinter, A. (1998). PARSER: A model for word segmentation. *Journal of Memory and Language, 39*, 246–263.
Roseberry, S., Richie, R., Hirsh-Pasek, K., Golinkoff, R., & Shipley, T. (2011). Babies catch a break: 7- to 9-month-olds track statistical probabilities in continuous dynamic events. *Psychological Science, 22*, 1422–1424.
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*, 27–52.
Saffran, J. R., Newport, E. L., Aslin, R. N., Tunick, R. A., & Barrueco, S. (1997). Incidental language learning: Listening (and learning) out of the corner of your ear. *Psychological Science, 8*, 101–105.
Stahl, A. E., Romberg, A. R., Roseberry, S., Golinkoff, R. M., & Hirsh-Pasek, K. (2014). Infants segment continuous events using transitional probabilities. *Child Development, 85*, 1821–1826.
Swingley, D. (2005). Statistical clustering and the contents of the infant vocabulary. *Cognitive Psychology, 50*, 86–132.
Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24*, 629–640.
Thiessen, E. D. (2010). Effects of visual information on adults' and infants' auditory statistical learning. *Cognitive Science, 34*, 1093–1106.
Thiessen, E. D. (2012). Effects of inter- and intra-modal redundancy on infants' rule learning. *Language Learning and Development, 8*, 197–214.

Toro, J. M., Sinnett, S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition, 97*, B25–B34.

Turk-Browne, N. B., Junge, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General, 134*, 552–564.

Yoshida, H., & Smith, L. B. (2005). Linguistic cues enhance the learning of perceptual cues. *Psychological Science, 16*, 90–95.