



ELSEVIER

Contents lists available at ScienceDirect

Journal of Experimental Child Psychology

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



Specificity of representations in infants' visual statistical learning



Dylan M. Antovich^{a,*}, Stephanie Chen-Wu Gluck^b, Elizabeth J. Goldman^c,
Katharine Graf Estes^a

^a University of California, Davis, Davis, CA 95616, USA

^b University of Oregon, Eugene, OR 97403, USA

^c University of California, Santa Cruz, Santa Cruz, CA 95064, USA

ARTICLE INFO

Article history:

Received 14 May 2018

Revised 21 November 2019

Available online 12 February 2020

Keywords:

Visual statistical learning

Object processing

Generalization

Infant cognition

Specificity of learning

Habituation paradigm

ABSTRACT

Past work has demonstrated infants' robust statistical learning across visual and auditory modalities. However, the specificity of representations produced via visual statistical learning has not been fully explored. The current study addressed this by investigating infants' abilities to identify previously learned object sequences when some object features (e.g., shape, face) aligned with prior learning and other features did not. Experiment 1 replicated past work demonstrating that infants can learn statistical regularities across sequentially presented objects and extended this finding to 16-month-olds. In Experiment 2, infants viewed test sequences in which one object feature (e.g., face) had been removed but the other feature (e.g., shape) was maintained, resulting in failure to identify familiar sequences. We further probed learning specificity by assessing infants' recognition of sequences when one feature was altered rather than removed (Experiment 3) and when one feature was uncorrelated with the original sequence structure (Experiment 4). In both cases, infants failed to identify sequences in which object features were not identical between learning and test. These findings suggest that infants are limited in their ability to generalize the statistical structure of an object sequence when the objects' features do not align between learning and test.

© 2019 Elsevier Inc. All rights reserved.

* Corresponding author.

E-mail address: dmantovich@ucdavis.edu (D.M. Antovich).

Introduction

As infants observe the world, they are confronted with a nearly continuous stream of visual information. Tracking regularities in this rich input helps infants to identify properties of objects (Wu, Gopnik, Richardson, & Kirkham, 2011), form categories (French, Mareschal, Mermillod, & Quinn, 2004; Sloutsky & Robinson, 2013), learn about the faces of their caregivers (Ferguson, Kulkofsky, Cashion, & Casasola, 2009; Gaither, Pauker, & Johnson, 2012), and begin to disambiguate the meanings of words (Smith & Yu, 2008), among many other feats. Given the ubiquity and significance of visual information surrounding an infant, it is not surprising that infants can track patterns in visual input in the form of conditional statistics, similar to those available in the auditory domain (Saffran, Aslin, & Newport, 1996). Past work has demonstrated that infants can detect the statistical structure of visual information in the form of contingencies between elements in a scene (Fiser & Aslin, 2002) and consistencies in the sequential presentation of objects (Bulf, Johnson, & Valenza, 2011; Kirkham, Slemmer, & Johnson, 2002; Kirkham, Slemmer, Richardson, & Johnson, 2007; Marcovitch & Lewkowicz, 2009; Slone & Johnson, 2015; Stahl, Romberg, Roseberry, Golinkoff, & Hirsh-Pasek, 2014; Tummeltshammer, Amso, French, & Kirkham, 2017). Although research has demonstrated that infants have impressive skills for detecting visual conditional statistics, there is still an incomplete understanding of the nature of this learning.

In an early test of infants' abilities to detect patterns in sequential visual information, Kirkham et al. (2002) presented 2-, 5-, and 8-month-olds with a stream of colorful shapes that formed consistent shape pairs, analogous to the bisyllabic words used in some investigations of statistical learning in the auditory domain (Antovich & Graf Estes, 2018; Graf Estes, 2012; Graf Estes, Gluck, & Bastos, 2015; Lew-Williams & Saffran, 2012). After habituating to a continuous object sequence, infants viewed sequences of object pairs that were consistent with exposure during learning or viewed random sequences of objects that violated the pair structure. Infants reliably discriminated between the random and structured sequences, indicating that they had learned the underlying statistical structure of the object sequence. Subsequent work has found that infants may possess the ability to track object co-occurrence patterns even from birth (see Bulf et al., 2011; reviewed in Krogh, Vlach, & Johnson, 2013), suggesting that infants may be particularly sensitive to this form of statistical information in the visual world.

Although evidence indicates that statistical regularities in the visual modality are accessible even to very young infants, little is known about the representations that underlie visual statistical learning. We do not yet know precisely what information infants encode and store about the elements within a sequence. There is recent evidence that infants store reliable visual sequences (i.e., a series of shapes) as cohesive units or "chunks." Slone and Johnson (2018) found that, after exposure to shape sequences, 8-month-olds differentiated sequences that occurred with high probability in the stream from sequences that partially overlapped with the high-probability sequences but did not form complete units. Their findings indicate that infants extract and store detailed information about the patterns of co-occurrence of the elements they observe and use this information to detect coherent units via statistical learning.

There is still a great deal to learn about the information infants encode, store, and apply from statistical learning experiences. The form of the representations that emerge from visual statistical learning is important because it may influence infants' processing of visual events in their environment. For instance, an infant may view a dinner routine in which plates, cups, and utensils are brought out and then food is served. Learning the properties of the set of objects involved may help the infant to understand and anticipate aspects of the routine. However, some object features may change across events (e.g., plates can be many colors and sizes, cups vary in height and material), whereas other features remain consistent (e.g., plates are mostly flat and round, cups are generally cylindrical). Thus, the ability to detect the set of features most relevant to the underlying sequence, and to ignore superfluous features, may affect infants' abilities to generalize learning and identify broader routines.

Work by Kirkham et al. (2007) suggests that infants can use multiple feature dimensions to encode and recognize consistent object sequences. Kirkham and colleagues examined spatiotemporal visual statistical learning during infancy, finding that 11-month-olds learned a sequence of locations for object pairs even when the objects were identical in color and shape. They also found that 8-month-olds required objects to have unique shapes and colors in order to detect changes in the sequence of object locations, whereas the youngest infants, 5-month-olds, only differentiated between familiar and violation sequences when all three features had changed. This suggests that infants as young as 5 months can use correlated features (e.g., color, shape, location) to learn the order of elements within a sequence and that older infants may require fewer redundant cues to identify these sequences. Similarly, Bertels, San Anton, Gebuis, and Destrebecqz (2017) found that infants between 8 and 12 months of age were able to use object location, color, and shape information to encode associations between the layout of a scene and the location of a target object. The work by Bertels et al. (2017) also suggests that infants relied on particular combinations of objects and locations to process the target–context pairings. This supports the role of local learning over global learning given that infants relied on only a portion of the object-based information available in a scene to identify regularities. The current work expands our understanding of the development of feature use in visual statistical learning by testing 16-month-olds' use of correlated features to learn and generalize regularities in a continuous stream of objects.

Related work by Turk-Browne, Isola, Scholl, and Treat (2008) has established that adult statistical learning functions over whole objects when features are correlated and functions over individual perceptual features when some features are irrelevant to learning the underlying object sequences. In the study, participants saw a stream of objects that differed from each other in both shape and color. Similar to Kirkham et al. (2002), objects were presented in consistent combinations (triplets in this case), with objects appearing on the screen one object at a time. At test, participants were asked to discriminate between consistent object triplets and random foils. Participants identified the previously seen sequences at greater than chance levels when the objects at test were identical to those seen during the learning phase. However, a series of experiments revealed that, despite initially learning the triplet structure, performance was diminished when participants viewed stimuli that removed some perceptual features of the objects at test (e.g., colorful shapes during training became monochromatic black shapes or identically shaped color swatches at test). Participants were also less successful at identifying target sequences when the objects' colors at test mismatched colors seen during the learning phase. This suggests that changing features between learning and testing presented a challenge. It inhibited adults' recognition of the trained sequence, but their learning was still sufficiently flexible to recognize the altered objects in their original sequences at levels that exceeded chance performance. Moreover, participants' recognition improved if color and shape failed to covary perfectly during learning. Specifically, when some objects appeared with a consistent shape but variable color during learning (the square was sometimes red, sometimes blue, and sometimes green), participants' recognition of the sequences was no longer disrupted by removing the color feature at test. This work suggests that adult visual statistical learning can occur over holistic objects or individual object features (e.g., shape), depending on how closely correlated the objects' features are. Adults' visual statistical learning adjusts to match the learning environment.

The current series of experiments investigated whether infant visual statistical learning can generalize to novel stimuli in which features have changed. Furthermore, we tested whether infant learning can adjust in response to variability in the input, similar to adults' flexible visual statistical learning (Turk-Browne et al., 2008). First, we replicated past work demonstrating that infants can track statistical regularities in visual sequences of shape- and color-based objects. Subsequent experiments examined whether infant statistical learning relies on precise object representations or on more abstract generalizable representations and whether the distinction is related to feature variability during learning.

The object processing literature suggests that infants use features such as shape, pattern, and color to individuate objects by the end of their first year of life (Wilcox & Chapa, 2004; Wilcox, 1999) and that infants form cohesive object representations based on visual feature co-occurrences (Barry, Graf Estes, & Rivera, 2015; Best, Yim, & Sloutsky, 2013; Fiser & Aslin, 2002; Wu et al., 2011). Thus, infants may track object co-occurrences by using correlated features to identify individual objects. If so,

infants' generalization of sequences acquired via visual statistical learning would be constrained by the presence of the co-occurring features. Alternatively, given that infants are able to detect consistent correlations among object features (Younger & Cohen, 1983; Younger, 1990), they may also be able to learn individual feature co-occurrences across objects (e.g., blue before red before green). If infants track multiple feature sequences simultaneously, this could allow them to generalize what they have learned to novel contexts even when some features present during learning are absent or changed at test. The current research addressed whether infants can generalize the structure of visual sequences to novel stimuli that share some, but not all, perceptual features with their original learning experience. Importantly, this differs from past research examining rule learning using visual sequences. For example, in work from Saffran, Pollak, Seibel, and Shkolnik (2007), 7-month-olds were presented with sequences of animals that followed ABB or ABA patterns. The authors found that infants could detect the ABB and ABA patterns when presented with many different sequences exemplifying a single pattern, prompting infants to acquire abstract rules that readily generalize across objects (see also Johnson et al., 2009). Infants in rule-learning tasks have also shown generalization across stimulus types (i.e., spoken ABB patterns recognized in nonspeech sounds; Marcus, Fernandes, & Johnson, 2007) following exposure to multiple examples of a pattern. On the other hand, the current work presented infants with consistent objects in consistent sequences, which may hinder, rather than promote, generalization.

Although infants' abilities to generalize statistical representations in the visual modality have not been assessed, analogous research in the auditory modality suggests that infants can generalize the output of auditory statistical learning to novel circumstances (e.g., new speaker voice; Graf Estes, 2012) and that increased variability during learning may enhance this flexibility (Graf Estes & Lew-Williams, 2015). Thus, it is possible that infants may recognize previously learned visual sequences even when features of these objects have been altered. However, given adults' difficulty in identifying previously viewed sequences when consistently co-occurring perceptual features changed between learning and recall (e.g., Turk-Browne et al., 2008), we predicted that infants would demonstrate statistical learning only when test conditions closely aligned with conditions present during learning.

Experiment 1

Experiment 1 was designed to replicate and build on past work demonstrating that infants readily track transitional probabilities in sequences of visual objects (e.g., Kirkham et al., 2002). During habituation, infants viewed a series of colorful shapes with a unique face superimposed on each object. Introducing faces to the colorful shapes allowed us to include a salient dimension to manipulate in Experiments 2–4. There is ample evidence that infants are highly attentive to faces (e.g., Frank, Vul, & Johnson, 2009); therefore, they are likely to attend to this feature along with the shape and color dimensions during learning. At test, we assessed whether infant looking time discriminated between sequences that were consistent with those learned during habituation and sequences that violated prior learning.

Method

Participants

A total of 22 16-month-old infants ($M = 16$ months 8 days, range = 15 months 12 days to 16 months 25 days) participated in Experiment 1. Previous research has focused on young infants (newborns: Bulf et al., 2011; 2–8-month-olds: Kirkham et al., 2002; Slone & Johnson, 2015), and there has been some recent research with older children (e.g., 5–12-year-olds: Bertels, Boursain, Destrebecqz, & Gaillard, 2015; Raviv & Arnon, 2018). The current study expanded this work to test older infants, who are presumably capable of more advanced visual pattern learning and generalization than very young infants (e.g., Kirkham et al., 2007; see Krogh et al., 2013, for discussion of the development of statistical learning abilities). Thus, selecting participants from an older age range than in prior work fills in a gap in our understanding of the developmental timeline of visual statistical learning and may

increase the likelihood of successful generalization in Experiments 2–4, where object features change between learning and test.

Across Experiments 1–4, infants had no history of known visual impairments and were born full-term. Infants primarily came from English-speaking households. In Experiment 1, seven infants heard a second language (up to 16 h per week), assessed via parent report. Additional infants were excluded from analyses because of fussiness or crying ($n = 10$) or moving out of the video frame ($n = 2$). For each experiment, before performing the full analyses, we examined the data for outliers based on looking-time difference scores (looking to violation vs. consistent test trials) greater than 2 *SD* from the mean. No infants met this criterion.

Stimuli

During the habituation phase, infants viewed silent videos presenting a series of 12 colorful shapes superimposed with cartoon faces (see Fig. 1). The objects appeared one by one on a white screen, loomed for 1 s until they doubled in size, reaching approximately 6×6 inches (9.6×9.6 degrees at a viewing distance of 36 inches) and then disappeared. The object sequence comprised repetitions of four consistent shape¹ triplets. These triplets were then pseudorandomized with the constraint that no object triplet could occur twice in a row. Thus, the transitional probability from one object to the next within each triplet was 1.00, and the average transitional probability across triplets was .33 (range = .26–.38). We created 21 videos for the habituation trials, each presenting a different randomization of the triplet sequences. The presentation order of the habituation videos was randomized.

After habituation, infants completed test trials to assess learning. Half of the test trials consisted of sequences that maintained the original object triplets, and half were randomized sequences. *Consistent sequences* had a statistical structure that aligned with the sequences viewed during habituation and included the same four object triplets presented in a pseudorandom order. There were four consistent sequence test trials with different pseudorandomized orders of the triplets, with the constraint that no triplet occurred twice in a row. *Violation sequences* were inconsistent with the order of objects seen during habituation. In these trials, the 12 individual objects were presented in a pseudorandom order, with the constraint that objects were not repeated twice in succession. Across the four violation sequence test trials, infants saw four different randomizations of the 12 shapes. The shapes for each experiment were created, animated, and edited using Microsoft PowerPoint, Adobe Photoshop, Adobe Director, and Final Cut Pro.

Procedure

Infants completed habituation and testing in a small, sound-attenuated booth. Throughout testing, each infant sat on a parent's lap facing a large television monitor used to display stimuli at infant eye level. Parents were discouraged from directing infants' attention while stimuli were playing and were monitored for interference. In addition, parents wore headphones playing classical music during the experiment to reduce the likelihood of infant–parent interaction. During habituation, infants viewed the object sequences described above. After each trial, infants saw an attention-getting video of a rotating pinwheel accompanied by music. The order in which the habituation videos were presented was randomized. Each habituation trial ended when an infant looked for a minimum of 2 s and looked away for at least 1 s or after 20 s had elapsed. The habituation phase ended once infants' average looking time across three consecutive trials fell below 50% of their average looking time during the first three habituation trials or after a total of 21 trials. Test trials began immediately after habituation trials. The eight test trials (four consistent and four violation test sequences) were presented in random order. If infants acquired the statistical structure of the visual stream during habituation, we expected them to attend longer to the test sequences that violated prior learning than to sequences that aligned with it. The experiment was conducted using Habit 2 software (Oakes, Sperka, & Cantrell, 2015). Experimenters coded infant looking behavior online using keypresses while watching a video feed of the testing booth from a separate room. Experimenters were unaware of trial type.

¹ We use the term “shape” to refer to the two-dimensional geometric shapes and their associated colors. These features co-occurred consistently.

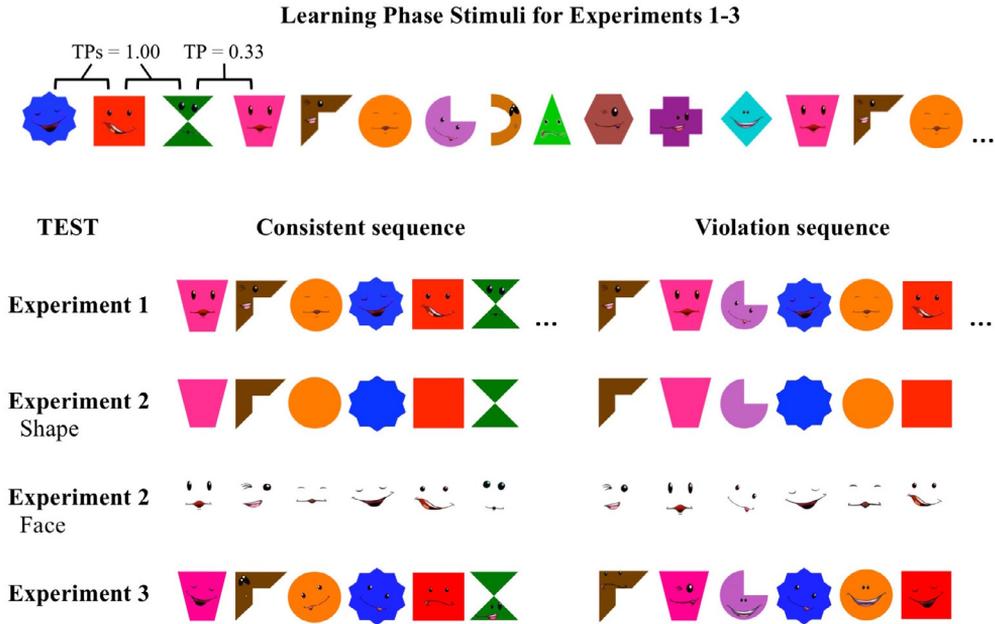


Fig. 1. Example learning and test phase stimuli for Experiments 1–3. Experiment 4 learning phase stimuli were similar to the test stimuli from Experiment 3, and Experiment 4 test stimuli were similar to the test stimuli from the Experiment 2 Shape condition. TP, transitional probability.

Results and discussion

Table 1 shows infants' time and number of trials to reach the habituation criterion for Experiments 1 to 4. As reported in Table 1, a multivariate analysis of variance (ANOVA) revealed that there were no significant differences in habituation time or number of trials to reach habituation across experiments, $F(4, 107) = 0.36, p = .83, \eta_p^2 = .01$.

A two-tailed, paired-samples *t* test comparing infants' average looking time toward consistent and violation sequences revealed that infants looked longer during the violation sequences, $t(21) = -2.61, p = .016, d_z = -0.55$ (see Table 2). To illustrate the range of infants' performance, Fig. 2 presents looking-time data for each infant using the looking-time difference scores (looking time on violation–consistent test trials). This finding shows that 16-month-olds detect statistical regularities present in visual sequences, in line with work by Kirkham et al. (2002). With infants' learning pattern for the object sequences established, subsequent experiments investigated infants' abilities to recall previously learned object sequences under changed circumstances.

To support interpretation of the findings, we also calculated Bayes factors (Rouder, Speckman, Sun, Morey, & Iverson, 2009; Wagenmakers et al., 2018). In contrast to null hypothesis testing methods, Bayes factors allow for a comparison of the strength of support for the null hypothesis and the alternative hypothesis. In other words, we can quantify evidence in favor of the absence of an effect (i.e., the null hypothesis) and evidence in favor of the presence of an effect (i.e., the alternative hypothesis). For all experiments, we used the calculator by Rouder et al. (2009) and the recommended JZS Bayes factor value, which is not based on prior expectations of effect sizes. The estimated Bayes factor indicated that the hypothesis that infants' looking times to violation versus consistent sequences were *different* was 3.3 times more likely than the null hypothesis of no difference, a substantial effect (Jeffreys, 1961).

Table 1

Average time (in seconds) and number of trials to habituation by experiment.

Experiment	Habituation time (SD)	Habituation trials (SD)
1	203 (73)	15.8 (5.0)
2 (Face)	238 (106)	16.0 (4.9)
2 (Shape)	235 (104)	15.8 (5.1)
3	192 (90)	14.9 (5.1)
4	195 (106)	14.5 (6.0)

Note. There were no significant differences in habituation time or trials across experiments (all $ps > .10$).

Table 2

Average looking times (in seconds) by experiment.

Experiment	Consistent trial looking (SD)	Violation trial looking (SD)
1	8.23 (3.19)	10.28 (4.22)
2 (Face)	11.37 (5.28)	11.34 (5.27)
2 (Shape)	10.03 (4.03)	10.36 (3.53)
3	10.04 (4.41)	10.14 (4.03)
4	9.94 (3.91)	10.21 (4.10)

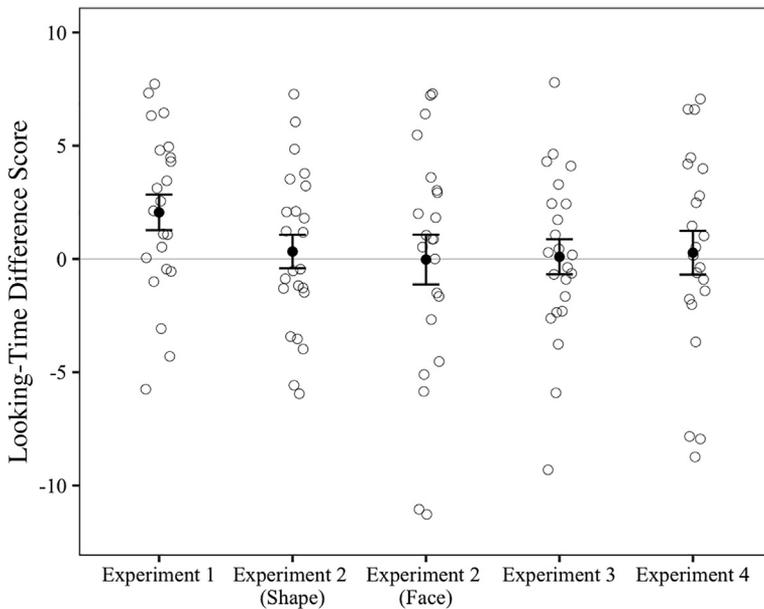


Fig. 2. Looking-time difference scores (average looking time on violation trials–looking time on consistent trials) for each infant by experiment. Filled points indicate the mean within each experiment, with error bars representing standard errors of the mean. Note that scores above zero indicate novelty preferences, whereas scores below zero indicate familiarity preferences.

Experiment 2

Experiments 2–4 built on the findings of Experiment 1 to explore the conditions under which infants recognize previously viewed sequences. We investigated whether infants can recognize familiar visual sequences when salient features of those sequences are removed or altered during testing. In

Experiment 2, we tested infants' recognition of the object sequences described in Experiment 1 when either the face or shape feature was removed. During learning, these features provided redundant cues to the statistical structure of the sequence. Here, we addressed whether infants require the presence of both these correlated cues at test to recognize the structures they learned.

Method

Participants

A total of 45 16-month-old infants ($M = 16$ months 17 days, range = 16 months 3 days to 17 months 0 days) participated in Experiment 2. Infants came from the same population described in Experiment 1. All infants primarily heard English, and 10 infants also heard a second language (up to 16 h per week), per parent report. Additional infants were excluded from analyses for fussiness or crying (Shape condition: $n = 13$; Face condition: $n = 12$), moving out of the frame (Face: $n = 1$), parental interference (Face: $n = 1$), falling asleep (Shape: $n = 1$), or equipment error (Shape: $n = 1$).

Stimuli

The object streams presented during habituation were identical to those from the habituation phase of Experiment 1. As in Experiment 1, at test infants viewed object sequences that were consistent with the statistical structure from habituation sequences (consistent trials) and sequences that violated prior learning (violation trials). Evidence from Experiment 1 suggests that infants can readily discriminate between these two trial types when all the objects' features are consistent between habituation and test. However, in Experiment 2, infants viewed test stimuli in which features of the objects had been removed from the stream. In the Shape condition, infants viewed test sequences in which the object shape and color were maintained but the object was no longer superimposed with a cartoon face (see Fig. 1). In the Face condition, test sequences retained only the faces from habituation without any background shape or color (see Fig. 1). Thus, during habituation the faces and shapes were perfectly correlated, as in Experiment 1, but at test salient features were removed. In both conditions, infants viewed consistent sequences that aligned with learning (but were missing a feature) and viewed violation sequences in which a feature was missing and the objects were presented in a novel randomized order.

Procedure

The procedure was the same as in Experiment 1, with the exception that during testing one group of infants viewed test sequences from the Shape condition ($n = 23$) and another group viewed test sequences from the Face condition ($n = 22$). If infants recognized previously seen sequences at test despite the removal of one of the objects' features, we expected infant looking time to differentiate between consistent and violation sequences. However, if removing a feature at test disrupted infants' recognition of the sequences, infant looking time should not differentiate between consistent and violation sequences.

Results and discussion

Table 1 shows infants' time and number of trials to reach habituation. We conducted a 2 (Condition: Face vs. Shape; between participants) \times 2 (Trial Type: consistent vs. violation; within participants) mixed ANOVA to determine whether infants' looking time discriminated between consistent and violation sequences at test and whether this discrimination differed across the two conditions. There were no significant main effects or interactions, suggesting that infants did not reliably discriminate between consistent and violation sequence test trials in either condition, nor did overall looking time differ across conditions, $F_s(1, 43) < 0.94$, $p_s > .34$, $\eta_p^2 < .02$ (see Table 2 and Fig. 2). It is possible that infants' abilities to discriminate between sequences of shapes or faces that aligned with or violated prior learning were impaired when a feature of the object was removed.

Although the interaction was not significant, given the visual salience of faces, we analyzed the data separately for the face-only and shape-only test conditions. In the Face condition, there was no difference in looking time to consistent versus violation sequences, $t(21) = 0.022$, $p = .983$,

$d_z = 0.004$. In the Shape condition, there was also no difference, $t(22) = 0.444$, $p = .662$, $d_z = 0.093$. The Bayes factor analysis was consistent with the t tests, providing further support for the null hypothesis. The Bayes factors indicated that the hypothesis that looking times did not differ was 4.5 times more likely than the alternative hypothesis of different looking times for the Face condition and was 4.18 times more likely for the Shape condition, both substantial effects (Jeffreys, 1961). Thus, the Bayes factor analyses support the conclusion that infants did not demonstrate recognition of the violation versus consistent test trials.

We also considered the possibility that, because all the test items were new (i.e., consistent and violation sequences had a missing dimension, face or shape), the novelty of the test items could have superseded infants' attention to the sequence-level novelty of the violation test sequences. To test this idea, we examined infants' performance separately for the initial and final sets of familiar and novel trials. We compared performance in the first two consistent trials versus first two violation trials and then compared the last two instances of each trial type. Because the trials were randomized, these analyses did not always split the first four versus last four test trials in sequence, but each comparison included an equivalent number of each trial type. We hypothesized that if the initial novelty of the altered stimuli overwhelmed infants' differentiation of the two trial types, discrimination may be apparent only in the latter portion of testing. Infants did not show a significant difference in looking time to the violation trials in the first set of trials, $t(44) = 0.332$, $p = .740$, $d_z = 0.049$, or in the second set of trials, $t(44) = -0.010$, $p = .992$, $d_z = -0.001$. Infants did not differentiate the violation and consistent test items as they became more familiar with them. We cannot rule out the possibility that the change in appearance of the test items affected performance. Rather, we propose that recognizing the consistent sequence across change is an important part of generalization and that infants have not displayed successful generalization under these conditions.

The use of faces in the current stimuli may have affected infants' learning. Faces are known to have privileged status for infants (Frank et al., 2009; Valenza, Simion, Cassia, & Umiltà, 1996). Although statistical learning is thought to be implicit, attentional resources are still necessary to capture the appropriate input (Toro, Sinnett, & Soto-Faraco, 2005). The attention-grabbing presence (and subsequent conspicuous absence) of faces could have affected how infants processed the shape sequence during learning and testing. Infants may have primarily attended to the sequence of faces, ignoring information provided by the background shapes, which could inhibit recognition of the shape sequences at test. However, as described above, we did not observe learning differences between the Shape and Face conditions of Experiment 2 or habituation differences between experiments. Thus, it is unlikely that the manipulation of face features specifically (in contrast to another salient feature, shape) can fully account for the observed effects.

During learning, the manipulated features (i.e., face and shape) covaried perfectly, with each shape being superimposed with a consistent and unique cartoon face. At test, the contingency between shape and face could not be used to support recognition of previously viewed sequences. Work on object processing has suggested that across the first year of life infants are increasingly able to use features such as shape, pattern, and color to individuate objects (Wilcox & Chapa, 2004; Wilcox, 1999) and that infants form holistic object representations via co-occurring features (Wu et al., 2011). Thus, when infants learn sequences in which object features are correlated, it may be difficult for infants to recognize the same statistical regularities when only a subset of object features is present at test.

It is also possible that infants' recognition of the object sequence structure was disrupted by the visual disparity between objects presented during learning versus at test. Seeing shapes with no faces and faces with no shapes may have prevented infants from generalizing their prior experience with the object sequences. If the test items were more holistically similar to learning (i.e., colorful shapes with superimposed faces), infants may recognize the consistent sequences even when test items are not identical to those viewed during learning (i.e., face-shape pairings differ between learning and test).

Experiment 3

Experiment 3 investigated the specificity of infants' statistical representations by assessing whether infants would generalize the statistical structure of previously viewed object sequences to

novel but visually similar stimuli. The structure of the task was comparable to that in Experiment 2; infants viewed the same sequences of shapes paired with consistent faces during the learning phase. During testing, infants viewed test trials that had shapes with superimposed faces, but the face–shape combinations were pseudorandomized. That is, at test the shape–face correlations were inconsistent with learning. With this manipulation, the overall appearance of the test items was similar to learning, but the precise feature combinations of the original objects were absent. Salient differences in the color and shape of individual objects may provide enough information for infants to recognize previously viewed sequences when they are coupled with similar, although not identical, face information. This allowed us to test whether infants can generalize visual statistical learning across a change in object feature combinations. If infants learn the sequence regularities separately for individual features (i.e., shape and face), their prior experience could support recognition of the previously viewed shape sequences even when the face feature is no longer informative.

Method

Participants

A total of 23 16-month-old infants ($M = 16$ months 16 days, range = 16 months 0 days to 16 months 28 days) participated in Experiment 3. Infants came from the same population described in Experiment 1. All infants primarily heard English, and six infants also heard a second language (up to 16 h per week), per parent report. Additional infants were excluded from analyses for fussiness or crying ($n = 17$), equipment error ($n = 2$), or distraction ($n = 1$).

Stimuli

The object streams presented during the habituation were identical to those in Experiment 1. At test, infants viewed streams of objects in which the shape sequences were either consistent with the habituation sequences (consistent trials) or violated past learning (violation trials). However, in Experiment 3, the shape–face contingencies that were present during learning were eliminated at test. The face sequence was pseudorandomized so that each shape appeared with several different faces. The constraints on the face pseudorandomization were that no face occurred twice in succession and that some shape–face pairs were not possible due to the shapes' dimensions. Each shape appeared with at least three different faces (see Fig. 1). Thus, the test objects were broadly visually similar to objects viewed during habituation (i.e., they consisted of shapes with superimposed faces) but test sequence objects used novel shape–face pairs.

Procedure

The procedure was identical to that in Experiment 1.

Results and discussion

Infants' number of trials and time to habituation are shown in Table 1. A two-tailed, paired-samples t test comparing infants' average looking time to consistent versus violation test sequences revealed that infants did not reliably discriminate between trial types, $t(22) = -0.12$, $p = .90$, $d_z = -0.02$ (see Table 2 and Fig. 2). Thus, infants provided no evidence that they recognized the previously encountered object triplets at test when the pattern presented by one feature of the objects (faces) mismatched the information encountered during learning. The Bayes factor analysis indicated that the hypothesis that looking times did not differ was 4.5 times more likely than the alternative hypothesis that looking times differed. This analysis supports the notion that infants treated the violation and consistent test items similarly.

As in Experiment 2, we also examined infants' looking time in the first two consistent trials versus the first two violation trials and did the same for the final consistent and violation trials. In the first set of trials, infants looked marginally longer to the consistent sequence than to the violation sequence, $t(22) = -1.89$, $p = .072$, $d_z = -0.394$. In the second set of trials, infants looked marginally longer to the violation sequence, $t(22) = 1.77$, $p = .090$, $d_z = 0.369$. These patterns did not reach the threshold for significance, so we must be cautious in their interpretation. There are several reasons why infants' pat-

tern of preference may have shifted over the course of testing. It is possible that in the early trials infants detected the similarity of the consistent shape sequence or the mismatch of the shape sequence with the randomized faces, which initially drew their attention. By the end of the test trials, attention shifted to the novelty of the violation sequences. Alternatively, infants may have used exposure to the consistent test sequences to gather additional experience with the shape sequences that were encountered during learning, promoting a shift in attention toward the novelty of the violation sequences by the end of testing. A third possibility is that because both the violation and consistent test trials included face sequence violations, infants treated the test trial types as equivalent (although the Face condition of Experiment 2 suggests that this might not be the case). It is also possible that the looking pattern represents random changes in attention over the course of eight test trials. These are interesting hypotheses that require additional experiments to pull apart and are beyond the scope of the current investigation. However, what we have demonstrated is that “filling in” the shapes with a different sequence of faces did not yield robust generalization, which would be demonstrated by longer looking to the violation test sequence, as in Experiment 1.

Thus, both Experiments 2 and 3 suggest that infants' abilities to generalize statistical relationships in the visual modality are limited, at least in the current paradigm. In comparison with learning demonstrated in Experiment 1, infants' overall looking patterns in Experiments 2 and 3 suggest that infants had difficulty in detecting the consistent statistical structure shared across learning and test sequences or were unable to express their learning when features of the objects were altered. Based on prior work with adults (Turk-Browne et al., 2008), we hypothesized that the consistent face–shape pairings presented during the learning phase led infants to develop precise representations of the objects, which then impaired subsequent performance at test when the objects were changed (Oakes, Coppage, & Dingel, 1997). In other words, the invariable feature pairings seen during learning may have prevented infants from forming more abstract generalizable representations. In this case, infants' generalization may improve if the object sequences presented during learning exhibited greater feature variability, highlighting sequence-relevant and -irrelevant features. Experiment 4 investigated this possibility.

Experiment 4

Experiment 4 explored whether increased feature variability during learning would allow infants to generalize previously encountered statistical regularities. Infants viewed streams of shape triplets that had superimposed faces, but the faces did not covary with shape during learning or present a reliable pattern on their own. We then tested whether infants could recognize the shape sequences without faces (as in Experiment 2) after training with objects that maintained consistent shape triplets but had variable faces. Thus, the faces were uninformative for learning or recognizing the shape sequences.

Method

Participants

A total of 22 16-month-old infants ($M = 16$ months 19 days, range = 16 months 0 days to 17 months 0 days) participated in Experiment 4. Infants came from the same population described in Experiment 1. Most infants primarily heard English, with nine infants also exposed to a second language (up to 16 h per week), per parent report. One additional infant was primarily exposed to Spanish. The pattern of results was the same with this infant excluded, so the infant was included in the final sample. Additional infants were excluded from analyses for fussiness or crying ($n = 11$), equipment error ($n = 2$), or excessive movement ($n = 1$).

Stimuli

The object streams presented during the habituation phase maintained the same shape triplet sequences from Experiments 1 to 3. However, the shapes were no longer paired with consistent superimposed faces. Each token of an individual shape was pseudorandomly paired with a face during

habituation using the same constraints as the test phase in Experiment 3 (see Fig. 1). Each shape appeared with a minimum of three faces. The face and shape features of the objects were not correlated during habituation, and the faces occurred in a pseudorandom order. Thus, the face information was uninformative for learning the shape sequences, and the faces did not form consistent sequences themselves. At test, infants viewed streams of shapes with a statistical structure that either was consistent with the habituation sequences or violated past learning; the test items were the same as the Shape condition in Experiment 2 and had no faces.

Procedure

The procedure was identical to that in Experiment 1.

Results and discussion

Infants' number of trials and time to habituation are shown in Table 1. A two-tailed, paired-samples *t* test comparing infants' average looking time to consistent and violation test trials revealed that infants did not reliably discriminate between trial types, $t(21) = -0.29$, $p = .90$, $d_z = -0.06$ (see Table 2 and Fig. 2). The Bayes factor analyses indicated that the hypothesis that looking times did not differ was 4.3 times more likely than the alternative hypothesis of different looking times. This analysis supports the notion that infants treated the violation and consistent test items similarly. Infants displayed no evidence of recognizing the shape sequence, suggesting that variability in the face–shape pairings did not help infants to recognize the shape-only sequence.

To test the possibility that the initial novelty of the test items overwhelmed infants' attention, we analyzed the first set of consistent and violation trials separately from the last set of trials. The results were the same in both sets; infants did not differentiate violation and consistent trials in the first set, $t(21) = 0.230$, $p = .752$, $d_z = 0.049$, or second set, $t(21) = 0.144$, $p = .887$, $d_z = 0.031$, of test trials.

This contrasts with the findings from experiments with adults demonstrating that variation in feature combinations during learning allowed adults to track statistical regularities across features rather than whole objects (Turk-Browne et al., 2008). The findings in the current experiment align with Experiments 2 and 3, suggesting that infant visual statistical learning does not readily generalize across changes in objects' appearances. However, from this experiment, it is not possible to determine whether infants failed to learn at all when the face and shape cues were uncorrelated. Future experiments will test whether variation in features impedes learning overall or whether infants trained with consistent shape sequences but variable faces can recognize the shape sequences when variable faces are also present at test. Regardless, Experiment 4 demonstrates that uncoupling the face–shape pairings does not help infants to recognize the shape sequences in isolation.

General discussion

The findings from this series of experiments suggest that infants' representations of statistically defined visual sequences are tightly tied to the perceptual details of the input. In Experiment 1, we found that infants learned a visual sequence of colorful shapes with superimposed faces. This result replicates prior demonstrations of robust visual statistical learning of sequences (Kirkham et al., 2002; Marcovitch & Lewkowicz, 2009; Slone & Johnson, 2015) using an older age group and stimuli with more feature dimensions than many previous experiments. Experiment 1 helps to bridge research on early visual statistical learning with the literature investigating visual statistical learning in older children and adults (e.g., Bertels et al., 2015; Raviv & Arnon, 2018; Turk-Browne et al., 2008; Turk-Browne & Scholl, 2009). However, the absence of evidence of learning in Experiments 2–4 suggests that there are limitations to infant visual statistical learning.

When some visual features of the objects changed between learning and test, infants did not appear to recognize the underlying statistical structure that was shared with habituation sequences; they no longer differentiated between consistent and violation test sequences. This failure to generalize from the original habituation sequences occurred when one feature was absent at test (i.e., face or shape; Experiment 2) and when both features were present at test but only one aligned with prior

learning (Experiment 3). Importantly, the stimuli presented during learning sequences in Experiments 1–3 were identical, suggesting that infants' difficulty in differentiating the consistent versus violation test sequences in Experiments 2 and 3 was due to the presence of altered object features at test rather than to an underlying failure to learn. Experiment 4 assessed whether additional variability during habituation would promote generalization in statistical learning by providing one relevant feature with a consistent statistical structure across objects (i.e., shape) and one irrelevant feature without this structure (i.e., face). We hypothesized that tracking statistical regularities for the relevant feature would allow infants to generalize learning to sequences in which the irrelevant feature was absent, as was found for adults (Turk-Browne et al., 2008). Contrary to this hypothesis, Experiment 4 demonstrated that training on shape sequences with variable face pairings during learning did not support the formation of more robust generalizable representations of the shape sequences. Across experiments, infants demonstrated what they had learned about the visual sequences only when the learning and test items matched on all dimensions. Infants' failure to discriminate consistent versus violation sequences in Experiments 2–4 cannot tell us whether infants are incapable of generalizing any form of sequential information acquired via visual statistical learning. However, it does suggest that generalization provides additional challenges and, thus, may require additional support (e.g., several redundant cues). This aligns with findings from Kirkham et al. (2007) suggesting that multiple correlated features can help infants to detect a change in pattern. It is possible that with scaffolding infants may have successfully discriminated between consistent and violation sequences in Experiments 2–4.

One alternative explanation for infants' lack of clear preference in the looking tasks in Experiments 2 and 4 is that the novel appearance of the test items (e.g., via removal of face or shape features) might have overwhelmed infants' attention, causing looking time for both trial types to be at "ceiling" for these stimuli. Thus, a novelty preference for violation test sequences may have been negated by a global increase in infants' attention to test trials due to the change in form between training and test. However, this alternative explanation seems unlikely given that infant performance was similar in the initial and final sets of test trials. The results of Experiment 3 also support this conclusion given that the test trials maintained the overall visual appearance of the objects (colorful shapes superimposed with a variety of cartoon faces), yet infants' looking times still did not reliably differentiate between trial types. Taken together, these findings suggest that the similarity of infants' looking time to consistent and violation sequences was not due to overwhelming novelty of both trial types at test.

Another alternative explanation was introduced in work by Addyman and Mareschal (2013), who found that infants could discriminate between structured and unstructured visual sequences even without prior training. Thus, it is possible that infants in Experiment 1 were simply differentiating between trial types based on local variability. However, if this were the case, we would anticipate similar discrimination at test in Experiments 2–4 because test trials in these experiments had similar structures to trials in Experiment 1 (i.e., ordered triplet sequences vs. randomized object sequences). This suggests that test performance in Experiment 1 was likely driven by learning during training rather than the inherently structured or unstructured nature of the test stimuli. To further test the alternative explanation, we had a group of infants from the same population described in Experiment 1 complete a control condition in which they viewed test trials identical to those in Experiment 1 but without prior training. Infants ($n = 22$) did not look reliably longer to consistent sequences (i.e., structured: $M = 13.99$ s, $SD = 3.99$) compared with violation sequences (i.e., unstructured: $M = 13.76$ s, $SD = 4.32$), $t(21) = 0.300$, $p = .767$, $d_z = 0.064$. This result suggests that infants in Experiment 1 had learned the underlying structure of the sequences during training and relied on this knowledge to discriminate between trial types at test.

Infants' difficulty in recognizing the altered visual sequences in Experiments 2 to 4 aligns with the extraction and integration framework of statistical learning (Erickson & Thiessen, 2015; Thiessen, Kronstein, & Hufnagle, 2013) and the notion that infants extract coherent chunks via visual statistical learning (Slone & Johnson, 2018). The extraction and integration framework proposes that chunks of consistently co-occurring elements are formed through an iterative process in which memory for consistent sequences strengthens and memory for inconsistent sequences decays due to time or interference. As chunked sequences are reencountered, the alignment of perceptual features between previously encountered and new stimuli matters; if the novel input differs from prior learning, recog-

dition is impaired due to a weakened memory trace. In the current study, infants successfully recognized previously learned sequences when perceptual features aligned perfectly between learning and test (e.g., Experiment 1) but did not reliably demonstrate recognition of the learned sequences when perceptual features changed between the two experimental phases (Experiments 2–4). This framework suggests that providing fewer correlated features and increasing noise during the learning phase of Experiment 4 may have increased the difficulty of the task, reducing the perceptual features available to support learning. If so, additional training with the variable face stimuli in Experiment 4 may help infants to ignore the uninformative variability in face features to detect regularities in the underlying shape triplets.

Differences in training and testing format limit our ability to compare the current infant findings with prior adult findings. Infant testing methods are often limited in their ability to detect fine-grained differences in learning, typically yielding the presence or absence of evidence of learning (i.e., statistically reliable differentiation of two test trial types) rather than graded differences in performance. In contrast, measures of adults' responses can reveal graded differences in accuracy even across conditions that show successful (i.e., above-chance) learning. Thus, more sensitive infant testing methods may allow us to better detect generalization. In any case, some limitations on infant learning in the current study were not shared by adults in work from Turk-Browne et al. (2008), suggesting that visual statistical learning is not a static ability but rather one that may change over time. Turk-Browne et al. reported that, for adults, changing the color or shape of items in a sequence moderated participants' performance relative to baseline learning. Like adults, 16-month-olds' expression of learning was affected by a salient feature change. However, infant and adult learners responded differently to variation during learning. For adults, when the features (color and shape) are not tightly coupled during learning, adults can shift to detecting patterns in isolated features (e.g., monochromatic shapes) rather than whole objects. In contrast, infants' generalization did not benefit from variation in feature combinations during learning. To further assess this pattern of learning, future work could examine whether infants learn the underlying structure of the object sequences when one feature dimension is uncorrelated and uninformative of the sequence's structure both at learning and at test.

Our results cannot directly address the hypothesis that visual statistical learning abilities develop over time. Nevertheless, differences between infants' performance in the current study and related adult work (Turk-Browne et al., 2008) align with growing evidence that visual statistical learning matures throughout infancy and childhood. Kirkham et al. (2007) test of spatiotemporal visual statistical learning showed that older infants could use a single object feature, location, to identify temporal sequences that aligned with prior learning, whereas younger infants required additional supporting perceptual features (i.e., shape and color) to learn the sequences. This suggests that infants' use of perceptual features in visual statistical learning develops across the first year of life. In addition, Raviv and Arnon (2018) found that performance in a visual statistical learning task improved steadily across middle childhood, demonstrating that visual statistical learning abilities continue to change well beyond infancy.

The current finding that infants failed to generalize in the visual statistical learning task differs from past work on linguistic statistical learning, which has demonstrated that infants can generalize statistical knowledge across changing contexts. Infants can recognize statistically segmented words in a novel voice (Graf Estes, 2012) and can detect the underlying statistical structure of speech presented across multiple voices (Graf Estes & Lew-Williams, 2015). It is possible that infants' frequent experience in navigating surface-level changes in speech input (e.g., changes in pitch, volume, and background noise) may have prepared infants for similar changes instituted in these statistical learning tasks. The visual objects in the current experiments are much less familiar than speech; infants have not had rich experience in determining how to interpret the objects or their features (see also Saffran et al., 2007). In this case, infants' attention to detail in the task may be an adaptive learning strategy. The highly novel context of tracking sequences of colorful shapes with faces may elicit a representation that precisely codes all the features of the stimuli. The specificity of this representation may inhibit recognition when the learning probe (feature-altered test items) does not effectively link back to the original highly specified memory. It would be possible to systematically test the specificity with which infants encoded the stimuli using a change detection task similar to Ross-Sheehy, Oakes, and

Luck (2003). In this case, infants would be presented with visual streams in which object features are altered or remain the same. If infants encoded the object representations with a high level of specificity, they should prefer the changing stream over the unchanging stream. Alternatively, if infants' memories for the objects did not include all discriminating features, they should not distinguish between the two streams.

Our findings suggest that infant visual statistical learning may, at least initially, be constrained by the specific examples infants encounter. From very early in life, infants can generalize patterns acquired via visual learning (Ferguson, Franconeri, & Waxman, 2018; Johnson et al., 2009; Saffran et al., 2007). In the current tasks, however, infants are primed to learn a stimulus-specific pattern rather than a general rule, and learning is accomplished over a very short time. Afterward, they face the challenge of recognizing a fragment of a familiar pattern with information missing or even contradicting prior learning. This fundamental difference in the nature of the exposure sequences may drive differences between the current work and prior studies examining rule learning. Moreover, the abstract nature of the stimuli may further hinder generalization. A recent meta-analysis from Rabagliati, Ferguson, and Lew-Williams (2019) suggests that pattern abstraction is easiest for stimuli that are meaningful, with patterns in speech detected more readily than nonspeech patterns. Given the novel nature of the current stimuli for infants, additional experience with the objects may have been necessary for infants to successfully generalize learning after changes to key features. Thus, as infants gain experience with series of objects and events in their environments, their abilities to recognize familiar patterns despite surface-level changes may become more robust. To continue an example introduced above, a change in the collection of objects or the sequence of events related to mealtime may be initially difficult to interpret for an infant, but greater familiarity with the sequence and experience with variability in the order and composition of events may promote generalization and support recognition despite surface-level changes.

In conclusion, our findings provide additional evidence that infants can track co-occurrence regularities across visual elements presented sequentially. However, infants' representations of the visual elements appear to be highly specified, leading missing or changed features to impair infants' abilities to recognize previously learned statistical sequences. Of course, the ability to generalize emerges with experience and becomes more robust across development; generalization is essential for applying learning in one situation and extending it to novel events. Future research is needed to investigate which types of experiences (e.g., greater exposure, greater variation in exposure) support successful generalization.

Declaration of Competing Interest

The authors declared that there is no conflict of interest.

Acknowledgments

This research was supported by a grant to Katharine Graf Estes from the National Institute of Child Health and Human Development (Grant No. HD062755). We thank parents who generously contributed their time, as well as Jessica Bloom, Rebekah Solomon, Destiny Harriman, Carolina Bastos, Lauren Levenson, Daniela Guarnizo, and the members of the Language Learning Lab at UC Davis for their assistance with this research. Thanks also go to Lauren Slone for helpful comments on an earlier draft of the manuscript.

Appendix A. Supplementary material

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jecp.2019.104772>.

References

- Addyman, C., & Mareschal, D. (2013). Local redundancy governs infants' spontaneous orienting to visual-temporal sequences. *Child Development, 84*, 1137–1144.
- Antovich, D. M., & Graf Estes, K. (2018). Learning across languages: Bilingual experience supports dual language statistical word segmentation. *Developmental Science, 21*, e12548.
- Barry, R. A., Graf Estes, K., & Rivera, S. M. (2015). Domain general learning: Infants use social and non-social cues when learning object statistics. *Frontiers in Psychology, 6*. <https://doi.org/10.3389/fpsyg.2015.00551>.
- Bertels, J., Boursain, E., Destrebecqz, A., & Gaillard, V. (2015). Visual statistical learning in children and young adults: How implicit?. *Frontiers in Psychology, 5*. <https://doi.org/10.3389/fpsyg.2014.01541>.
- Bertels, J., San Anton, E., Gebuis, T., & Destrebecqz, A. (2017). Learning the association between a context and a target location in infancy. *Developmental Science, 20*, e12397.
- Best, C. A., Yim, H., & Sloutsky, V. M. (2013). The cost of selective attention in category learning: Developmental differences between adults and infants. *Journal of Experimental Child Psychology, 116*, 105–119.
- Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition, 121*, 127–132.
- Erickson, L. C., & Thiessen, E. D. (2015). Statistical learning of language: Theory, validity, and predictions of a statistical learning account of language acquisition. *Developmental Review, 37*, 66–108.
- Ferguson, B., Franconeri, S. L., & Waxman, S. R. (2018). Very young infants learn abstract rules in the visual modality. *PLoS One, 13*(1), e190185.
- Ferguson, K. T., Kulkofsky, S., Cashion, C. H., & Casasola, M. (2009). The development of specialized processing of own-race faces in infancy. *Infancy, 14*, 263–284.
- 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., Vul, E., & Johnson, S. P. (2009). Development of infants' attention to faces during the first year. *Cognition, 110*, 160–170.
- French, R. M., Mareschal, D., Mermillod, M., & Quinn, P. C. (2004). The role of bottom-up processing in perceptual categorization by 3- to 4-month-old infants: Simulations and data. *Journal of Experimental Psychology: General, 133*, 382–397.
- Gaither, S. E., Pauker, K., & Johnson, S. P. (2012). Biracial and monoracial infant own-race face perception: An eye tracking study. *Developmental Science, 15*, 775–782.
- Graf Estes, K. (2012). Infants generalize representations of statistically segmented words. *Frontiers in Psychology, 3*. <https://doi.org/10.3389/fpsyg.2012.00447>.
- Graf Estes, K., Gluck, S. C.-W., & Bastos, C. (2015). Flexibility in statistical word segmentation: Finding words in foreign speech. *Language Learning and Development, 11*, 252–269.
- Graf Estes, K., & Lew-Williams, C. (2015). Listening through voices: Infant statistical word segmentation across multiple speakers. *Developmental Psychology, 51*, 1517–1528.
- Jeffreys, H. (1961). *The theory of probability* (3rd ed.). New York: Oxford University Press.
- Johnson, S. P., Fernandes, K. J., Frank, M. C., Kirkham, N., Marcus, G., Rabagliati, H., & Slemmer, J. A. (2009). Abstract rule learning for visual sequences in 8- and 11-month-olds. *Infancy, 14*, 2–18.
- 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 Psychology, 3*. <https://doi.org/10.3389/fpsyg.2012.00598>.
- Lew-Williams, C., & Saffran, J. R. (2012). All words are not created equal: Expectations about word length guide infant statistical learning. *Cognition, 122*, 241–246.
- Marcovitch, S., & Lewkowicz, D. J. (2009). Sequence learning in infancy: The independent contributions of conditional probability and pair frequency information. *Developmental Science, 12*, 1020–1025.
- Marcus, G. F., Fernandes, K. J., & Johnson, S. P. (2007). Infant rule learning facilitated by speech. *Psychological Science, 18*, 387–391.
- Oakes, L. M., Coppage, D. J., & Dingel, A. (1997). By land or by sea: The role of perceptual similarity in infants' categorization of animals. *Developmental Psychology, 33*, 396–407.
- Oakes, L. M., Sperka, D. J., & Cantrell, L. (2015). *Habit 2*. University of California, Davis, Center for Mind and Brain (unpublished software).
- Rabagliati, H., Ferguson, B., & Lew-Williams, C. (2019). The profile of abstract rule learning in infancy: Meta-analytic and experimental evidence. *Developmental Science, 22*, e12704.
- Raviv, L., & Arnon, I. (2018). The developmental trajectory of children's auditory and visual statistical learning abilities: Modality-based differences in the effect of age. *Developmental Science, 21*, e12593.
- Ross-Sheehy, S., Oakes, L. M., & Luck, S. J. (2003). The development of visual short-term memory capacity in infants. *Child Development, 74*, 1807–1822.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review, 16*, 225–237.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*, 1926–1928.
- Saffran, J. R., Pollak, S. D., Seibel, R. L., & Shkolnik, A. (2007). Dog is a dog is a dog: Infant rule learning is not specific to language. *Cognition, 105*, 669–680.
- Slone, L. K., & Johnson, S. P. (2015). Infants' statistical learning: 2- and 5-month-olds' segmentation of continuous visual sequences. *Journal of Experimental Child Psychology, 133*, 47–56.
- Slone, L. K., & Johnson, S. P. (2018). When learning goes beyond statistics: Infants represent visual sequences in terms of chunks. *Cognition, 178*, 92–102.

- Sloutsky, V. M., & Robinson, C. W. (2013). Redundancy matters: Flexible learning of multiple contingencies in infants. *Cognition*, *126*, 156–164.
- Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, *106*, 1558–1568.
- 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.
- Thiessen, E. D., Kronstein, A. T., & Hufnagle, D. G. (2013). The extraction and integration framework: A two-process account of statistical learning. *Psychological Bulletin*, *139*, 792–814.
- Toro, J. M., Sinnett, S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition*, *97*, B25–B34.
- Tummeltshammer, K., Amso, D., French, R. M., & Kirkham, N. Z. (2017). Across space and time: Infants learn from backward and forward visual statistics. *Developmental Science*, *20*, e12474.
- Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 399–407.
- Turk-Browne, N. B., & Scholl, B. J. (2009). Flexible visual statistical learning: Transfer across space and time. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 195–202.
- Valenza, E., Simion, F., Cassia, V. M., & Umiltà, C. (1996). Face preference at birth. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 892–903.
- Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., ... Morey, R. D. (2018). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin & Review*, *25*, 35–57.
- Wilcox, T. (1999). Object individuation: Infants' use of shape, size, pattern, and color. *Cognition*, *72*, 125–166.
- Wilcox, T., & Chapa, C. (2004). Priming infants to attend to color and pattern information in an individuation task. *Cognition*, *90*, 265–302.
- Wu, R., Gopnik, A., Richardson, D. C., & Kirkham, N. Z. (2011). Infants learn about objects from statistics and people. *Developmental Psychology*, *47*, 1220–1229.
- Younger, B. A. (1990). Infants' detection of correlations among feature categories. *Child Development*, *61*, 614–620.
- Younger, B. A., & Cohen, L. B. (1983). Infant perception of correlations among attributes. *Child Development*, *54*, 858–867.