

Infants Segment Continuous Events Using Transitional Probabilities

Aimee E. Stahl
Johns Hopkins University

Alexa R. Romberg
Indiana University

Sarah Roseberry
University of Washington

Roberta Michnick Golinkoff
University of Delaware

Kathryn Hirsh-Pasek
Temple University

Throughout their 1st year, infants adeptly detect statistical structure in their environment. However, little is known about whether statistical learning is a primary mechanism for event segmentation. This study directly tests whether statistical learning alone is sufficient to segment continuous events. Twenty-eight 7- to 9-month-old infants viewed a sequence of continuous actions performed by a novel agent in which there were no transitional movements that could have constrained the possible upcoming actions. At test, infants distinguished statistically intact units from less predictable ones. The ability to segment events using statistical structure may help infants discover other cues to event boundaries, such as intentions, and carve up the world of continuous motion in meaningful ways.

Imagine a frequent event in the life of an infant, like bath time. A parent might place the infant in the tub, open a bottle and put soap in their hands, wash the infant, and rinse. These actions are likely to occur in the same order each time this event happens. On a daily basis, infants observe and engage in routines like this that are composed of reliable sequences of actions. The event of bath time may be followed by other events that also have predictable structure, like putting on pajamas or book reading. How do infants know when one event ends and another begins? Segmenting events into units is critical for many skills, including anticipating future actions, imitating others, categorizing events, and learning words that label those actions. An outstanding question for developmental scientists is how infants parse the action sequences that make up events in a way that scaffolds these skills.

One hypothesis for how event perception progresses is that infants begin with basic, domain-general learning mechanisms that allow them to group actions based on the sequential predictability of the actions they observe (Baldwin & Baird, 2001; Baldwin, Baird, Saylor, & Clark, 2001; Roseberry, Richie, Hirsh-Pasek, Golinkoff, & Shipley, 2011). Infants could use these initial groupings to discover more abstract cues to event structure, such as the actor's intentions, which are known to play a role in adults' global event segmentation (e.g., Wilder, 1978; Zacks, 2004; Zacks & Tversky, 2001). A similar arc has been proposed for word segmentation. Infants initially use basic perceptual learning to parse the speech stream, in which they perceive highly predictable sequences of syllables as more word-like than less predictable sequences (e.g., Aslin, Saffran, & Newport, 1998; Graf Estes, Evans, Alibali, & Saffran, 2007; Romberg & Saffran, 2010; Saffran, Aslin, & Newport, 1996). Infants then generalize across familiar words to find more language-specific cues to word boundaries, such as lexical stress (Sahni, Seidenberg, & Saffran, 2010; Thiessen, Kronstein, & Hufnagle, 2013; Thiessen & Saffran, 2003).

This study was supported by an NSF Graduate Research Fellowship to Aimee E. Stahl, NIH Ruth Kirschstein NRSA training Grant T32 HD007475-17 to Indiana University, and NICHD Grant RO1HD050199 and NSF Grant 0642632 to Roberta Michnick Golinkoff and Kathryn Hirsh-Pasek. We thank Sujeet Ranganathan and Katherine Ridge for their help with data collection, Emily Tynan for her assistance with creating the stimuli, and all of the parents and infants who participated.

Correspondence concerning this article should be addressed to Aimee E. Stahl, Department of Psychological & Brain Sciences, Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218. Electronic mail may be sent to astahl4@jhu.edu.

© 2014 The Authors
Child Development © 2014 Society for Research in Child Development, Inc.
All rights reserved. 0009-3920/2014/8505-0006
DOI: 10.1111/cdev.12247

Several studies have demonstrated infants' visual sequence learning skills by testing whether they can track regularities in static features like shape and color (Bulf, Johnson, & Valenza, 2011; Kirkham, Slemmer, & Johnson, 2002). However, event sequences consist of dynamically changing movements rather than static features. Eight-month-old infants are sensitive to the sequential statistics of actions performed by a human agent (Roseberry et al., 2011). Infants viewed a sequence of hand motions in which some motions reliably followed other motions, forming units that were combined into larger sequences. Similar to real-world actions, unique transitional movements connected each hand motion to the one that came after. For example, transitioning from the motion of pressing palms together to the motion of forming an "X" with one's arms requires that one hand pass in front of the other. This motion is markedly different than transitioning from pressing palms to stacking one's fists. These unique transitions provided cues to the sequential structure much like coarticulation in fluent speech, where the production of one sound is influenced by the pronunciation of the preceding or following sound. In Roseberry et al. (2011), the next action was constrained to those physically compatible with the trajectory of the transitional motion from the preceding action. Thus, infants had two cues to the sequential structure: sequential regularities and the transitional movement that connected one motion to the next.

There is no doubt that such physical constraints are present in everyday actions and likely provide useful information about upcoming motions. However, unique transitions are not required for auditory sequence learning with either words (e.g., Gómez, 2002; Lany & Gómez, 2008) or tones (e.g., Saffran, Johnson, Aslin, & Newport, 1999). Therefore, this study sought to test whether infants can segment action sequences based solely on their sequential predictability without the additional cue of transitional movements that physically constrain the upcoming motion.

A finding that infants can segment actions based only on sequential predictability would bolster the claim that statistical learning operates similarly across domains. Conversely, a finding that infants must rely on physical constraints as well as sequential structure would raise the possibility that statistical learning of action sequences relies on a different mechanism than statistical learning in other domains. Although researchers have concluded that finding predictable structure in sequential stimuli is a fundamental learning mechanism (Kirkham et al., 2002; Krogh, Vlach, & Johnson, 2013), it has yet to

be established that the same learning mechanism accounts for all types of statistical learning (Conway & Christiansen, 2006; Thiessen et al., 2013).

One recent study found that *adults* segmented a continuous action sequence based on sequential predictability alone (Baldwin, Andersson, Saffran, & Meyer, 2008). The materials in that study consisted of videos of an actor performing simple actions. The actor started and finished each action in the same position, thereby eliminating additional cues about the upcoming action. We applied this design technique to ask whether infants also can segment actions based solely on sequential predictability. In addition, we tested whether infants can segment actions performed by a completely novel (nonhuman) agent or whether segmentation relies on the familiarity of the stimulus (e.g., hands, such as in Roseberry et al., 2011).

In this study, infants were shown an animated agent (a starfish with eyes) performing whole-body actions (e.g., spin, twist, bend). Each action started and ended in a fully extended "star" position so that actions could flow naturally and continuously without transitional movements providing additional cues to the sequential structure of the units. Actions were arranged into 4 three-action units that formed the full, continuous sequence. If infants' ability to track sequential structure in actions relies on physically constraining transitional motions between actions, or on stimulus familiarity, they should fail to segment the sequence and therefore not recognize the three-action units at test. If infants possess robust perceptual learning skills for action segmentation, as they do for language, infants should discriminate between the highly predictable three-action units and sequences that are less predictable.

Method

Participants

Participants were 28 monolingual, English-reared infants (13 female) between the ages of 7 and 9 months ($M = 236$ days, $SD = 18.4$, range = 209–266 days). One additional infant was excluded for being distracted by external noise. Familiarization data from 4 of the 28 infants were not saved due to a computer error. The final sample was predominantly Caucasian and middle class.

Stimuli and Procedure

Infants sat on a parent's lap facing the monitor displaying the stimulus. Parents were instructed to

close their eyes and refrain from speaking. An observer, blind to the stimulus sequence viewed by the infant, watched the infant on another monitor and indicated by key press when the infant was looking at the screen. Two computers controlled the stimuli through Habit X software (Cohen, Atkinson, & Chaput, 2004): One controlled the auditory attention-getter during familiarization; the other controlled visual stimuli during familiarization and audio and visual presentations during the test phase.

Familiarization

Stimuli consisted of a purple animated starfish with eyes in the center of a black screen. The starfish performed 12 actions (bow, clap, backflip, turn, jumping jack, back bend, side bend, twist, arm circles, split, toe touch, scrunch; see the Appendix). Each action was a 1-s cycle that began and ended in a fully extended "star" position. The return to the "star" position allowed actions to be combined in any order and ensured that transitional movements between actions did not provide cues to the sequence structure.

The 12 actions were organized into 4 unique three-action units (Figure 1). Within-unit transitional probabilities (TPs) were 1.0. The event corpus consisted of a total of 42 units. Two of the three-action units occurred 14 times (high-frequency units) and the other two occurred 7 times (low-frequency units). The units were arranged in the sequence to create "part-units" that spanned the last action of one high-frequency unit and the first two actions of the other high-frequency unit. These part-units had an internal TP of 0.5 between the first and second actions, and occurred the same number of times as the low-frequency action units (Aslin et al., 1998). Infants viewed one of two familiarization corpora counterbalanced for which action units were high and low frequency.

The familiarization video played continuously for its 126-s duration. If infants looked away from

the screen for 1 s, an auditory attention-getter (a futuristic noise) played until they resumed looking at the screen. The video continued to play during the auditory attention-getter.

Test

Immediately after familiarization, infants viewed a series of 12 test trials, presented in three blocks of four test sequences. Two of these sequences were the low-frequency units and the other two were the part-units described above. Because the low-frequency units and the part-units occurred with equal frequency in the familiarization corpus, successful discrimination required that infants recognize the differing TP structure between these items. The low-frequency intact units in one corpus were the part-units in the other, and vice versa, which ensured that infants did not simply prefer particular action sequences over others during the test trials.

The test trials were presented in one of four pre-set orders that varied which test item was presented first; the same type of trial never appeared more than twice in a row. Each test trial was preceded by an attention-getter (a looming geometric pattern accompanied by a futuristic sound). The test trial began once the infant looked at the screen. Test sequences were looped until infants looked away for 1 s or for a maximum duration of 15 s.

Results

Infants were generally attentive during the experiment. The average looking time to the familiarization stimulus was 94.9 s ($SD = 17.2$ s, range = 57.1–119.8 s) of the 126-s video. Because each individual action was 1 s long, the minimum time required to distinguish the test items was 1.5 s, which was enough to see the first action and the beginning of the second (the crucial pair to distinguish intact from

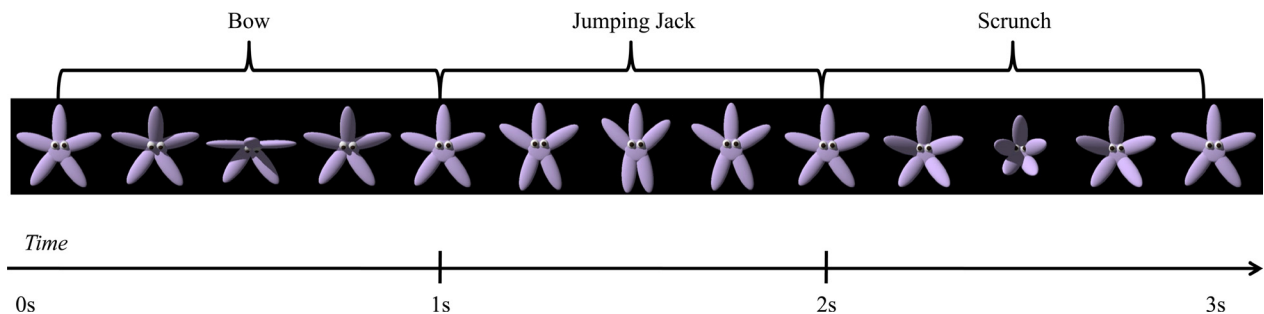


Figure 1. Sample sequence of actions from the familiarization corpus. These static images represent what was shown in the dynamic events.

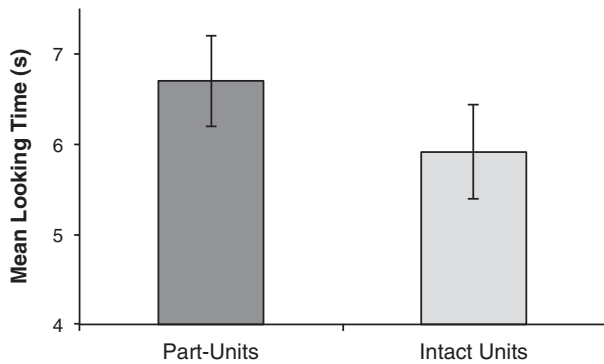


Figure 2. Mean looking times at test for part-units and intact units. Bars represent standard error of the mean.

part-units). Test trials on which infants looked for less than 1.5 s were excluded. This led to the exclusion of 12 (3.5%) trials. All infants contributed at least 4 test trials of each trial type to the analysis.

A preliminary analysis revealed no significant differences in looking times for the different counterbalanced corpora ($p > .18$) and so all analyses were collapsed across corpora. If infants used TPs to segment the three-action units from the continuous stream, then part-units with a mid-sequence dip in TP should be novel compared to intact units. As predicted, infants looked longer to part-units ($M = 6.70$ s, $SD = 2.68$ s) than to intact units ($M = 5.91$ s, $SD = 2.75$ s), and this difference was statistically significant, $t(27) = 2.09$, $p = .046$. Group means are shown in Figure 2; 19 (68%) of 28 infants looked longer to the part-units than to the intact units (two-sided binomial test, $p = .087$).

Because infants were free to look away from the screen during the familiarization phase, they varied in how much time they spent viewing the action sequence. Familiarization time was positively correlated with *total* looking time during test, indicating that some infants were more attentive overall, $r(22) = 0.45$, $p = .03$. However, the amount of time infants looked during familiarization was not correlated with their looking time difference to intact and part-units, $r(22) = 0.08$, $p = .68$. Thus, it was not the case that infants who looked more during familiarization showed stronger novelty preferences at test. Infants' age and gender were also unrelated to their difference scores ($ps > .6$).

Discussion

This study sought to determine whether infants can use TPs alone to segment events, and our results demonstrate infants' keen sensitivity to statistical

structure in action sequences. Infants successfully segmented a continuous stream of actions performed by a novel agent, looking significantly longer to part-units than to intact units. Importantly, the part-units and intact units used to test segmentation occurred with equal frequency during familiarization. Thus, infants must have utilized the probability that one action followed another to discriminate between intact and part-units.

Several studies have demonstrated infants' sensitivity to the distribution of static features in visual sequences (e.g., Fiser & Aslin, 2002; Kirkham et al., 2002). This study crucially extends this body of work by revealing that infants track sequential regularities in continuously changing actions without predictive transitional movements between them (see also Roseberry et al., 2011). Our results provide additional evidence that statistical learning is a fundamental learning mechanism that operates over many domains in similar ways, and that tracking distributional information may play an important role in event segmentation in infancy. However, given that not all infants looked longer to part-units than intact units, it is likely that there are individual differences in infants' statistical learning abilities that contribute to their event segmentation.

These findings add to a growing body of literature investigating the skills underlying infant event segmentation and suggest that perceptual learning may provide the foundation for learning higher level cues for action processing, such as the agent's intention (e.g., Baldwin et al., 2001; Saylor, Baldwin, Baird, & LaBounty, 2007). Before infants have top-down knowledge of intentions, they may begin to carve up events in the world based on sequential predictability. The observation of reliable sequences of actions has been implicated in sequence learning by nonhuman species. After repeated viewings of action sequences, young gorillas can parse individual actions and extract their sequential predictability without necessarily understanding the intention of the action sequence. This "string parsing" may ultimately help them to produce the sequence of behaviors themselves (Byrne, 1999). Infants in our experiment faced a similar challenge, as there were no cues to causes, goals, or intentions. Rather, they had to extract the statistical regularities only by observing the sequence of actions.

The ability to segment a continuous stream of actions enables infants to perceive events as consisting of coherent units. This may make processing more efficient by allowing them to predict downstream events and form hierarchical event representations (e.g., first comes bath time, then comes

book reading, both of which are part of the larger event of “getting ready for bed”). Event segmentation may even enhance memory for those events (Zacks, Speer, Vettel, & Jacoby, 2006; Zacks & Swallow, 2007). Furthermore, in order for infants to themselves produce action sequences (e.g., bathing themselves or a baby doll), they need to recognize and parse the individual components (e.g., wash then rinse). Tracking sequential predictability might be one means by which they achieve comprehension and then production of meaningful action sequences.

Several open questions remain to be addressed in future research. First, how does sequential predictability influence infants’ ability to learn a label for a novel action sequence? Infants’ formation of novel label–object mappings is facilitated when labels consist of predictable syllable sequences based on immediately prior linguistic exposure (Graf Estes et al., 2007; Hay, Pelucchi, Graf Estes, & Saffran, 2011). Likewise, infants may more readily map labels to novel events with higher internal consistency based on their prior experience. Second, future studies should examine whether and when infants generalize links created by distributional information from one agent to another. Once infants map labels to such sequences, they will need to extend them to new contexts to use language generatively (e.g., the action sequence of bath time can be labeled as such whether dad or the babysitter performs it). Finally, how does the presence of other cues, like intentions (Baldwin et al., 2001), familiarity (Hespos, Saylor, & Grossman, 2009), or goals (Loucks & Meltzoff, 2013), influence the role of sequential predictability in event segmentation? In natural language processing, word segmentation is facilitated by the conjunction of multiple cues, such as sequential predictability, familiar words, and isolated utterances (Lew-Williams, Pelucchi, & Saffran, 2011; Mersad & Nazzi, 2012). Tracking distributional information in action sequences likely contributes to event processing, but event segmentation will no doubt be most robust when multiple cues align.

This study suggests that during the 1st year of life infants possess the ability to segment events using sequential predictability. This domain-general mechanism may support the development of domain-specific and higher order processes, a vital precursor to carving up the world in meaningful ways.

References












- Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, *9*, 321–324. doi:10.1111/1467-9280.00063
- Baldwin, D., Andersson, A., Saffran, J. R., & Meyer, M. (2008). Segmenting dynamic human action via statistical structure. *Cognition*, *106*, 1382–1407. doi:10.1016/j.cognition.2007.07.005
- Baldwin, D. A., & Baird, J. A. (2001). Discerning intentions in dynamic human action. *Trends in Cognitive Sciences*, *5*, 171–178. doi:10.1016/S1364-6613(00)01615-6
- Baldwin, D. A., Baird, J. A., Saylor, M. M., & Clark, M. A. (2001). Infants parse dynamic action. *Child Development*, *72*, 708–717. doi:10.1111/1467-8624.00310
- Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition*, *121*, 127–132. doi:10.1016/j.cognition.2011.06.010
- Byrne, R. W. (1999). Imitation without intentionality. Using string parsing to copy the organization of behaviour. *Animal Cognition*, *2*, 63–72. doi:10.1007/s100710050025
- Cohen, L. B., Atkinson, D. J., & Chaput, H. H. (2004). *Habit X: A new program for obtaining and organizing data in infant perception and cognition studies (Version 1.0)*. Austin: University of Texas.
- Conway, C. M., & Christiansen, M. H. (2006). Statistical learning within and between modalities: Pitting abstract against stimulus-specific representations. *Psychological Science*, *17*, 905–912. doi:10.1111/j.1467-9280.2006.01801.x
- 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. doi:10.1073/pnas.232472899
- Gómez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science*, *13*, 431–436. doi:10.1111/1467-9280.00476
- Graf Estes, K. M., Evans, J., Alibali, M. W., & Saffran, J. R. (2007). Can infants map meaning to newly segmented words? Statistical segmentation and word learning. *Psychological Science*, *18*, 254–260. doi:10.1111/j.1467-9280.2007.01885.x
- Hay, J. F., Pelucchi, B., Graf Estes, K., & Saffran, J. R. (2011). Linking sounds to meaning: Infant statistical learning in a natural language. *Cognitive Psychology*, *63*, 93–106. doi:10.1016/j.cogpsych.2011.06.002
- Hespos, S. J., Saylor, M., & Grossman, S. (2009). Infants’ ability to parse continuous actions. *Developmental Psychology*, *45*, 575–585. doi:10.1037/a0014145
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence of a domain general learning mechanism. *Cognition*, *83*, B35–B42. doi:10.1016/S0010-0277(02)00004-5
- Krogh, L., Vlach, H. A., & Johnson, S. P. (2013). Statistical learning across development: Flexible yet constrained. *Frontiers in Language Sciences*, *3*, 598. doi:10.3389/fpsyg.2012.00598
- Lany, J., & Gómez, R. L. (2008). Twelve-month-olds benefit from prior experience in statistical learning. *Psychological Science*, *19*, 1247–1252. doi:10.1111/j.1467-9280.2008.02233.x

- Lew-Williams, C., Pelucchi, B., & Saffran, J. R. (2011). Isolated words enhance statistical language learning in infancy. *Developmental Science, 14*, 1323–1329. doi:10.1111/j.1467-7687.2011.01079.x
- Loucks, J., & Meltzoff, A. N. (2013). Goals influence memory and imitation for dynamic human action in 36-month-old children. *Scandinavian Journal of Psychology, 54*, 41–50. doi:10.1111/sjop.12004
- Mersad, K., & Nazzi, T. (2012). When Mommy comes to the rescue of statistics: Infants combine top-down and bottom-up cues to segment speech. *Language Learning and Development, 8*, 303–315. doi:10.1080/15475441.2011.609106
- Romberg, A. R., & Saffran, J. R. (2010). Statistical learning and language acquisition. *Wiley Interdisciplinary Reviews: Cognitive Science, 1*, 906–914. doi:10.1002/wcs.78
- Roseberry, S., Richie, R., Hirsh-Pasek, K., Golinkoff, R. M., & Shipley, T. (2011). Babies catch a break: Seven- to 9-month-olds track statistical probabilities in continuous dynamic events. *Psychological Science, 22*, 1422–1424. doi:10.1177/0956797611422074
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*, 1926–1928. doi:10.1126/science.274.5294.1926
- 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. doi:10.1016/S0010-0277(98)00075-4
- Sahni, S. D., Seidenberg, M., & Saffran, J. R. (2010). Connecting cues: Overlapping regularities support cue discovery in infancy. *Child Development, 81*, 727–736. doi:10.1111/j.1467-8624.2010.01430.x
- Saylor, M., Baldwin, D., Baird, J., & LaBounty, J. (2007). Infants' on-line segmentation of dynamic human action. *Journal of Cognition and Development, 8*, 113–128. doi:10.1207/s15327647jcd0801_6
- 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. doi:10.1037/a0030801
- Thiessen, E. D., & Saffran, J. R. (2003). When cues collide: Statistical and stress cues in infant word segmentation. *Developmental Psychology, 39*, 706–716. doi:10.1037/0012-1649.39.4.706
- Wilder, D. A. (1978). Effect of predictability on units of perception and attribution. *Personality and Social Psychology Bulletin, 4*, 281–284. doi:10.1177/014616727800400222
- Zacks, J. M. (2004). Using movement and intentions to understand simple events. *Cognitive Science, 28*, 979–1008. doi:10.1016/j.cogsci.2004.06.003
- Zacks, J. M., Speer, N. K., Vettel, J. M., & Jacoby, L. L. (2006). Event understanding and memory in healthy aging and dementia of the Alzheimer type. *Psychology & Aging, 21*, 466–482. doi:10.1037/0882-7974.21.3.466
- Zacks, J. M., & Swallow, K. M. (2007). Event segmentation. *Current Directions in Psychological Science, 16*, 80–84. doi:10.1111/j.1467-8721.2007.00480.x

- Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin, 127*, 3–21. doi:10.1037//0033-2909.127.1.3

Appendix

Static Depictions of Each Dynamic 1-S Action in the Familiarization Corpus

Bow	
Clap	
Backflip	
Turn	
Jumping Jack	
Back Bend	
Side Bend	
Twist	
Arm Circles	
Split	
Toe Touch	
Scrunch	