Cognition 146 (2016) 251–263

Contents lists available at ScienceDirect

Cognition

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

Infants use temporal regularities to chunk objects in memory

Melissa M. Kibbe [a,b,*], Lisa Feigenson [b]

[a] Department of Psychological and Brain Sciences, Boston University, 64 Cummington Mall, Boston, MA 02215, USA
[b] Department of Psychological and Brain Sciences, Johns Hopkins University, 3400 N Charles St, Baltimore, MD 21218, USA

ARTICLE INFO

Article history:
Received 20 February 2014
Revised 13 July 2015
Accepted 29 September 2015

Keywords:
Working memory
Chunking
Statistical learning
Infants

ABSTRACT

Infants, like adults, can maintain only a few items in working memory, but can overcome this limit by creating more efficient representations, or "chunks." Previous research shows that infants can form chunks using shared features or spatial proximity between objects. Here we asked whether infants also can create chunked representations using regularities that unfold over time. Thirteen-month old infants first were familiarized with four objects of different shapes and colors, presented in successive pairs. For some infants, the identities of objects in each pair varied randomly across familiarization (Experiment 1). For others, the objects within a pair always co-occurred, either in consistent relative spatial positions (Experiment 2a) or varying spatial positions (Experiment 2b). Following familiarization, infants saw all four objects hidden behind a screen and then saw the screen lifted to reveal either four objects or only three. Infants in Experiment 1, who had been familiarized with random object pairings, failed to look longer at the unexpected 3-object outcome; they showed the same inability to concurrently represent four objects as in other studies of infant working memory. In contrast, infants in Experiments 2a and 2b, who had been familiarized with regularly co-occurring pairs, looked longer at the unexpected outcome. These infants apparently used the co-occurrence between individual objects during familiarization to form chunked representations that were later deployed to track the objects as they were hidden at test. In Experiment 3, we confirmed that the familiarization affected infants' ability to remember the occluded objects rather than merely establishing longer-term memory for object pairs. Following familiarization to consistent pairs, infants who were not shown a hiding event (but merely saw the same test outcomes as in Experiments 2a and b) showed no preference for arrays of three versus four objects. Finally, in Experiments 4 and 5, we asked whether infants also remembered the specific identities of the objects in each chunk. In Experiment 4, we confirmed that infants remembered objects' identities in smaller arrays that did not require chunking. Next, in Experiment 5, we asked whether infants also remembered objects' identities in larger arrays that had been chunked on the basis of temporal regularities. Following a familiarization phase identical to that in Experiment 2a, we hid all four objects and then revealed either these same four objects, or four objects of which two had unexpectedly changed shape and color. Surprisingly, infants failed to look longer at the identity change outcome. Taken together, our results suggest that infants can use temporal regularities between objects to increase memory for objects' existence, but not necessarily for objects' identities.

1. Introduction

Research has revealed surprising limits on the amount of information that can be retained over brief intervals. Adults appear able to store representations of just three or four items at a time in visual working memory (e.g., Alvarez & Cavanagh, 2004; Cowan, 2001; Luck & Vogel, 1997; Sperling, 1960), and by around 10 months of age, infants show a similar memory limit across a range of experimental paradigms (e.g., Feigenson & Carey, 2003; Feigenson, Carey, & Hauser, 2002; Oakes, Hurley, Ross-Sheehy, & Luck, 2011; Ross-Sheehy, Oakes, & Luck, 2003; Zosh, Halberda, & Feigenson, 2011). For example, 12- to 21-month old infants who saw two or three objects hidden in a box, then saw just a subset of those objects retrieved, correctly searched the box for the missing object(s). In contrast, infants who saw four objects hidden and then saw any subset retrieved failed to keep searching (Barner, Thalwitz, Wood, & Carey, 2007; Feigenson & Carey, 2003; Feigenson & Carey, 2005; Feigenson & Halberda, 2004). This suggests that infants were unable to maintain a representation of four hidden objects, or even just a subset of the four. Hence working

* Corresponding author at: Department of Psychological and Brain Sciences, Boston University, 64 Cummington Mall, Boston, MA 02215, USA.
E-mail addresses: kibbe@bu.edu (M.M. Kibbe), feigenson@jhu.edu (L. Feigenson).

http://dx.doi.org/10.1016/j.cognition.2015.09.022
0010-0277/

memory in infants and young children appears to hold no more than about three items at a time.

Although working memory is constrained in adults, children, and infants, all of these populations have been shown to overcome these constraints through the use of chunking. In a chunked representation, individual items are grouped together but are still recoverable as individuals—this allows for the storage of more information in memory. For example, experienced chess players represent unified configurations of chess pieces on a game board (e.g., "Anastasia's Mate"), and can mentally "unpack" these higher-level representations into their constituent pieces. Adults can form these kinds of efficient, chunked representations using a variety of cues, including items' shared color or spatial proximity (Bower, 1972; Hitch, Burgess, Towse, & Culpin, 1996)—as well as more conceptual cues such as common category membership or semantic relatedness (Bower, Clark, Lesgold, & Winzenz, 1969; Chase & Simon, 1973; Ericsson, Chase, & Faloon, 1980; Mathy & Feldman, 2012). Recent work shows that chunking has its origins early in development. Fourteen-month old infants successfully remembered the presence of four hidden objects when the objects were presented in two spatially grouped sets of two before they were hidden, but not when these same objects were first presented in a single set of four (Feigenson & Halberda, 2004; Rosenberg & Feigenson, 2013). Like adults, 14-month-old infants can chunk using their knowledge of object categories: they remembered four total objects when an array contained two tokens of two different types (e.g., two cats and two cars), but not when the array contained four tokens of a single type (e.g., four different cats) (Feigenson & Halberda, 2008).

This research with young children shows that from early in development, working memory makes use of "snapshot" regularities. That is, when objects within an array share a common feature or spatial location that can be observed in a single glance, the objects can be represented more efficiently. However, snapshot regularities often are not available – for example, all of the objects in a scene may be unique, or may be evenly distributed in space. In such cases, is chunking possible?

A reason to suspect that it might be is that snapshot regularities are not the only source of information that may support chunking—other, more dynamic cues might also be used. For example, the frequency with which an object occurs in a local environment can be a powerful means to more efficient representation (e.g., Huffman, 1952). How reliably particular objects are seen together, the relative timing of objects' appearances, and objects' relative spatial positions are all temporal regularities that could potentially be used to create higher-order representations. Such temporal regularities are distinct from snapshot regularities in that they are unobservable from a single exposure. Instead, they must be gleaned from experiences that unfold over time. For example, an array of four evenly spaced, differently colored objects might contain no snapshot grouping cues. But if some smaller subset of the objects had previously been observed to occur together with high regularity, then this information, accumulated over time, might be useful as a basis for chunking.

Recent evidence suggests that adults can use these kinds of temporal regularities to increase the amount of information they remember from a visual scene. Brady, Konkle, and Alvarez (2009) showed adult observers eight simultaneously presented, differently colored circles for 1000 ms, after which the circles disappeared and adults were prompted to recall the color of one of them. The critical manipulation was whether, across trials, some of the colors were highly likely to appear next to each other. For some participants, particular colors often appeared together (e.g., a red circle appeared next to a blue circle on 80% of trials). For others, the color relationships were random across trials. Brady and colleagues found that within minutes, participants who observed the regularities outperformed participants who saw randomly configured arrays. They concluded that adults used the regularities that unfolded over time to compress item representations in memory, storing representations of pairs of co-occurring items more efficiently than they could store representations of two unrelated items.

Extracting regularities across events is a potentially powerful means of overcoming working memory limits in the absence of snapshot grouping cues. To achieve this, temporal regularities must be learned rapidly, and then used to encode an array with greater efficiency than if such regularities were absent or not yet learned. Previous studies show that young infants are indeed sensitive to temporal regularities in both visual and auditory stimuli (Aslin, Saffran, & Newport, 1998; Kirkham, Slemmer, & Johnson, 2002; Saffran, Aslin, & Newport, 1996a; Teinonen, Fellman, Näätänen, Alku, & Huotilainen, 2009; see Krogh, Vlach, & Johnson, 2013 for a review). Infants can parse artificial streams of continuous speech using conditional probabilities between syllables (Saffran, Newport, & Aslin, 1996b) and can parse visual streams of sequentially presented shapes using similar information (e.g., Kirkham et al., 2002). Nine-month old infants have been shown to use the co-occurrence statistics of visual elements to extract multi-part objects from scenes containing many smaller elements; they looked longer at pairs of elements previously seen to reliably co-occur than at pairs with a lower co-occurrence (Fiser & Aslin, 2002). However, it remains unknown whether these infants used co-occurrence statistics to form new, higher order memory representations that were available for further computation, or whether infants simply preferred looking at visual elements that had been statistically associated. One way to find out is to ask whether memory is more efficient when infants are provided with temporal chunking cues than when they are not—that is, whether experiencing temporal regularities among items increases the number of items infants can store in working memory. If so, it would suggest that working memory can be efficiently organized using a wide range of information types from early in development.

Here, in six experiments, we asked whether infants use regularities in object appearances over time in order to increase working memory performance. Because previous studies investigating preverbal chunking abilities examined infants of around 13–14 months old (Feigenson & Halberda, 2004; Feigenson & Halberda, 2008; Rosenberg & Feigenson, 2013), we focused on infants of a similar age. But unlike these previous studies, which presented infants with object arrays that could be chunked in a single glance, here we tested infants' memory for arrays of evenly spaced objects that each had a unique color and shape; thus any single viewing of an array contained no information that could be used for chunking.

In Experiment 1, we first confirmed the previously observed upper limit on the number of objects infants can remember from arrays lacking chunking cues. Following a familiarization phase, infants saw four unique objects hidden behind an occluding screen. The screen was then lifted to reveal either all four objects, or only three. We found that, as predicted, infants showed no visual preference between these two outcomes: they apparently failed to remember the presence of four hidden objects, even when each object had distinctive features, and even when they had been familiarized with all of the objects before the memory test. Next we asked whether infants would successfully remember the same array if first given the opportunity to experience temporal regularities among the objects. Infants in Experiment 2a were familiarized with the same four objects from Experiment 1, but this time saw the objects in successive presentations of pairs with multiple temporal regularities: object identity was yoked such that the appearance of one particular object (e.g., red disk) perfectly predicted the appearance of another (e.g., blue cross). In addition to this