UvA-DARE (Digital Academic Repository)

Distributional Information Shapes Infants' Categorization of Objects

Junge, C.; van Rooijen, R.; Raijmakers, M.

DOI

[10.1111/infa.12258](https://doi.org/10.1111/infa.12258)

Publication date

2018

Document Version

Final published version

Published in

Infancy

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Junge, C., van Rooijen, R., & Raijmakers, M. (2018). Distributional Information Shapes Infants' Categorization of Objects. *Infancy*, 23(6), 917-926. <https://doi.org/10.1111/infa.12258>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Distributional Information Shapes Infants' Categorization of Objects

Caroline Junge and Rianne van Rooijen

*Department of Experimental Psychology
Helmholtz Institute
Utrecht University*

Maartje Raijmakers

*Educational Studies
Faculty of Behavioural and Movement Sciences
Vrije Universiteit Amsterdam and
Department of Developmental Psychology and Amsterdam Brain & Cognition
University of Amsterdam*

While distributional learning has been successfully demonstrated for auditory categorization, this study tests whether this mechanism also applies to object categorization: Ten-month-olds ($n = 38$) were familiarized with either a unimodal or bimodal distribution of a visual continuum. Using automatic eye tracking, we assessed categorization through the alternating/nonalternating paradigm. For infants in the bimodal condition, their average dwell time was larger for alternating trials than for nonalternating trials, while infants in the unimodal condition initially looked equally long at both types of trials. This group difference suggests that the shape of frequency distribution bears on the number of categories that infants construct from a continuum. Later in test, all infants show this alternating preference. We conclude that categorization is a flexible process, continuously adjusting itself to additional input.

For infants, each day is filled with novel information. Being able to categorize proves a vital ability in reducing the diversity in the outside world, as it enables infants to put novel objects into familiar categories (Rosch, 1978). Such “perceptual categorization” (Mandler, 2000) is a prerequisite for word learning, as all exemplars within a basic level category share the same linguistic label (Waxman & Lidz, 2006). How do infants build such categories? Certainly, infants can be guided by top-down knowledge, that is, the linguistic label can shape categorization (e.g., Plunkett, Hu, & Cohen, 2008;

Waxman & Gelman, 2009). Yet it remains unclear on what kind of mechanisms infants rely when building novel concepts in the absence of language. This study examines whether frequency distributions of novel objects can guide infants' perceptual categorization.

Research shows that already three-month-olds can build perceptual categories, at an age before they acquire matching labels (e.g., Oakes & Ribar, 2005; Quinn, Eimas, & Rosenkrantz, 1993). One of the key mechanisms that infants rely on is perceptual similarity, which holds that entities that belong to the same category look more similar to one another than entities from different categories (Mandler, 2000; Medin & Coley, 1998; Rosch, 1978; Sloutsky, 2003). This weighing of perceptual similarity appears dynamic rather than fixed, as infants' attention to relevant dimensions shifts with experience and across different contexts (Smith & Heise, 1992).

Another mechanism fundamental for categorization is statistical learning, which is the ability to extract patterns across various stimuli (for a review, see Aslin & Newport, 2012). Most evidence that infants are sensitive to (various forms of) statistical regularities originates from speech perception research. For instance, one source of information is the probability of how elements in the input follow each other as it informs whether or not strings of elements belong together. Saffran, Aslin, & Newport (1996) were the first to demonstrate that infants can use transitional probabilities to find multisyllabic words in continuous speech. This type of statistical learning is not modality-specific but appears to be domain-general as infants exposed to a temporal sequence of visual shapes can recognize recurring patterns as well (Bulf, Johnson, & Valenza, 2011; Fiser & Aslin, 2002; Kirkham, Slemmer, & Johnson, 2002).

Keeping track of the distributional information along a continuum is another type of statistical learning as it enables the learner to detect whether or not there are multiple categories present. High-frequency tokens are generally considered to be more representative of a category than perceptually similar tokens with lower frequency (e.g., Nosofsky, 1988). There is evidence that object frequency can modulate the perception of perceptual similarity. Even when perceptual similarity remains constant, people rate those similar to high-frequency exemplars as higher in similarity than to low-frequency exemplars of a contrasting category (e.g., Nosofsky, 1988). Consequently, when there is little overlap between high-frequency tokens in a shared perceptual space, this indicates that they belong to distinct categories rather than one overarching category (Goudbeek & Swingley, 2006; Shepard, 1962).

Again, the first evidence that infants are sensitive to frequency distributions comes from speech perception literature. Maye, Werker, & Gerken (2002) familiarized infants with a few minutes of speech sounds from an 8-step [da]–[ta] continuum. Although all tokens were presented at least once during familiarization, some tokens were presented more often than others. Groups of infants differed in whether they listened to a sound distribution that reflected a single broad category ("unimodal distribution"; mainly midpoint exemplars 4 and 5) or to one that reflected two smaller categories ("bimodal distribution"; mainly endpoint exemplars 2 and 7). Crucially, the intervening tokens (3 and 6) were presented equally often to all infants. Following familiarization, both groups were tested with the Stimulus Alternation Preference ("SAP") paradigm (Best & Jones, 1998), in which infants listened to trials containing either alternations of tokens 3 and 6 ("alternating trials") or repetitions of one of these tokens ("nonalternating trials"). The SAP paradigm is based on the idea that infants generally prefer patterned stimuli over repetitive stimuli. Thus, a listening preference would

signal that infants detect that alternating trials differ from nonalternating trials. Results reveal that only infants in the two-peaked group differentiated between alternating and nonalternating trials, presumably because they mapped tokens 3 and 6 to distinct categories. In contrast, infants in the unimodal condition treated different tokens as equivalent to repetitions of same tokens, which is a hallmark of categorical perception (Goldstone & Hendrickson, 2010). This shows that infants' category formation can hinge on the shape of the frequency distribution presented during learning.

Since this seminal study (Maye et al., 2002), there have been several studies evidencing that the shape of frequency distributions guides categorization, yet all focused on (auditory) phoneme acquisition (e.g., Maye, Weiss, & Aslin, 2008; ter Schure, Junge, & Boersma, 2016; Yoshida, Pons, Maye, & Werker, 2010). It thus remains an open question whether infants employ this learning mechanism also outside the spoken language domain, that is, whether it serves as a domain-general cue.

Only one recent electrophysiological study has examined whether frequency distributions affect infants' discrimination of faces in a picture-repetition task (Altvater-Mackensen, Jessen, & Grossmann, 2017). During learning, all infants saw all exemplars of faces from a continuum, with one female gradually blending into another. At test, only infants familiarized with a bimodal distribution recognized whether or not a face was a direct repetition from the previous face they saw. This study suggests that the distributional learning mechanism is domain-general, as its scope extends to face recognition.

However, one could argue that face recognition might not be ideal for demonstrating that this distributional learning mechanism is domain-general. Face recognition requires fine discrimination skill (i.e., within-class distinctions), whereas categorization typically relies on between-class discriminations (i.e., creating in- and exclusive criteria; Mervis & Rosch, 1981; Oakes & Spalding, 1997). Moreover, infants appear specialized in face processing, like they are in language: compared to other types of object processing, face processing happens fast, automatic, and cannot be ignored (for a review see Johnson, Senju, & Tomalski, 2015). As a result, infants already have ample experience with viewing faces. A better test would be a situation in which infants have no experience with (a continuum of) visual objects. Will infants also use frequency distributions to build novel categories for objects they lack experience with?

There is some, indirect, evidence that infants indeed rely on distributional learning for object categorization. Given a continuum of cartoon drawings of animals, there are two studies that demonstrate that the composition during familiarization can guide categorization; however, in this case it is not the frequency distribution of all exemplars, but the magnitude of the correlation among features from a subset of exemplars that predicts categorization (Plunkett et al., 2008; Younger, 1985). Thus, these visual categorization studies do not directly address the key assumption of distributional learning, that is, that typical exemplars will be present more often in the environment than atypical exemplars.

This study therefore tests whether frequency distributions shape object categorization by adopting the experimental design from Maye et al. (2002). We familiarized infants with an 8-step continuum of two novel objects presented in a unimodal or bimodal distribution and subsequently tested infant discrimination with the SAP paradigm. Compared to the original design, we made two additional modifications.

First, we tested infants at 10 months rather than at 6–8 months, to match age groups of other studies on visual categorization (Plunkett et al., 2008; Younger, 1985), and because at this age, infants have become sensitive to feature correlations (Younger & Cohen, 1983).

Second, to give infants sufficient time to process the visual stimuli, we presented each token for 1,500 msec, while in the original paradigm, auditory tokens lasted only 465 msec. This entailed that while the familiarization phase maintained the original number of tokens and their frequencies (See Figure 1a), its length nearly tripled due to the nature of our stimuli.

For both groups of infants, testing was identical and involved alternations of end tokens 1 and 8 as well as direct repetitions of near-mid-tokens 3 or 6. Yet research suggests that emerging categories are continuously updated as sensitivity to relevant features is not innate but emerges with increased exposure (Schyns, Goldstone, & Thibaut, 1998). Hence, it is likely that the frequency of tokens presented during test will continue to shape infants' categorization. Figure 1 depicts the frequency distributions per group after familiarization (Figure 1a) and after testing (Figure 1b, with testing including additional presentations of only tokens 1, 3, 6, and 8). Although the frequency distribution for the bimodal group remains two-peaked after testing, the frequency distribution for the unimodal group appears less indicative of representing one large category. Indeed, studies on distributional learning often report a difference between groups only in the first test block (e.g., Feldman, Myers, White, Griffiths, & Morgan, 2013; ter Schure et al., 2016; Yeung & Nazzi, 2014). Therefore, we added test block as an additional within-subjects factor in our analyses. We hypothesized that if the frequency distribution of exemplars affects object categorization, infants with a bimodal distribution would prefer alternating over nonalternating trials, whereas infants who observed a unimodal distribution do not show any preference. If the hypothesized group effects appear short-lived, this would further indicate that infants' novel categories are fragile and can be adjusted with additional input.

METHODS

Participants

Via our subject pool, we recruited full-term healthy infants (mean age = 10.1, range 9.6–11.1 months) without prior exposure to our stimuli. We randomly assigned 38 infants to the unimodal ($n = 19$; 14 girls) or bimodal condition ($n = 19$; nine girls).

Figure 1 Distributions of visual stimuli. The X-axis represents the 8-step token continuum and the Y-axis token frequency, for both the unimodal (solid line) and bimodal (dotted line) conditions, aggregated after (a) familiarization phase; (b) familiarization plus test phase (testing consists of tokens 1, 3, 6, and 8).

Another 25 infants were excluded from analysis due to failing to calibrate the eye-tracker ($n = 5$); less than 20% total looking time during either familiarization ($n = 9$; range 0–17.7%) or test ($n = 7$, range 9.8–18.0%); or not contributing at least one test trial per block per trial type ($n = 4$). The study was approved by the ethical committee from the psychology department at the University of Amsterdam (2013-OP-2839) and conducted according to the declaration of Helsinki. All parents provided written informed consent, and their child received a small gift.

Stimuli

We selected two green unfamiliar cuddly toys from the Giant microbes series (the “nerve cell” and the “sea sparkle”; cf. www.giantmicrobes.com). Using Adobe Photoshop Elements 11, we first modified the source images to make them similar in height (around 730 pixels) against the same dark gray background ($1,280 \times 1,024$ pixels). The two toys were morphed using the Squirrel 2.1 software (Xiberpic.com) to create an eight-step continuum from 100% nerve cell to 100% sea sparkle (see Figure 2).

As in Maye et al. (2002), we used fillers to de-emphasize the experimentally relevant dimensions. We selected two other toys from the same series (the “cough” and the “meningitis”) and created four tokens from each filler toy by matching their colors to a similar shade of green as the target stimuli and slightly warping their shapes. In contrast to the target stimuli, these fillers do not form a continuum.

Design and procedure

Each infant sat in a car seat facing a 17" computer monitor at 60 cm distance in a semi-dark room. The parent sat next to the child and was instructed not to interfere. Infants' looking behavior was recorded with an EyeLink 1000 eye-tracker (Version 4.594; sample rate 500 Hz; Arm Mount Remote configuration). The experiment was programmed in SR Research Experiment Builder (SR Research Ltd., Mississauga, ON, Canada). The experiment commenced after a 4-point calibration. Its familiarization phase comprised four blocks of 24 stimuli each: 16 target stimuli from the continuum and eight different filler tokens. Per block, all infants saw all tokens from the continuum at least once (i.e., at least four times across blocks), but groups differed in which tokens they saw more frequently (i.e., maximum was 16 times across blocks; cf. Figure 1a). Within a block, the stimuli were presented at random. All stimuli were presented for 1,500 msec and accompanied by the same musical sound (a musical chimes). Blocks were separated by an attention getter (a looming ringing bell) to give infants a little break in the stream of images.

After familiarization, the experimenter had the option to recalibrate the eye-tracker before the test phase started automatically with the SAP paradigm (Best & Jones,

Figure 2 The 8-step continuum of the visual stimuli, ranging from the nerve cell (token 1) gradually morphing into the sea sparkle (token 8).

1998). There were eight test trials presented in interleaved order: four nonalternating (repetitions of near-mid-tokens 3 or 6) and four alternating trials (alternations between the endpoints 1 and 8). Trials lasted 22 sec, with tokens visible for 2 sec and an isi of 0.75 sec. Figure 3 illustrates both types. We counter-balanced order of test trials between subjects. Each test trial was followed by the ringing bell to regain the attention of the child.

RESULTS

Using the eye-tracker software (Dataviewer; SR Research Ltd.), we calculated for each child the total dwell time to the screen per trial. This measure represents per trial the summation of all fixations on screen and hence is most comparable to “total trial looking time” used in behavioral looking time experiments. Moreover, this measure is not as susceptible to large differences in eye-tracking data quality often observed in infant research as number of dwells or average dwell duration are (e.g., Hessels, Niehorster, Kemner, & Hooze, 2017). Across familiarization, groups did not differ in total looking time ($t(36) = 1.25$, $p = .22$). The unimodal group attended the screen for 101.3 sec ($SD = 25.7$), while the bimodal group did so for 89.9 sec ($SD = 30.4$).

To examine group differences during test, we conducted a 2 (Trial type: alternating vs. nonalternating) $\times$ 2 (Block: 1 vs. 2) repeated measures ANOVA, with Group (unimodal vs. bimodal) as a between-subject variable. We removed test trials with <500 msec looking time. Results from the ANOVA reveal two main effects. Across groups, infants’ looks decreased over blocks (e.g., main effect of Block ($F(1, 36) = 17.2$, $p < .001$, $\eta_p^2 = .32$); block 1: $M = 12.0$ sec, $SD = 4.55$; block 2: $M = 9.57$, $SD = 4.42$), and infants generally prefer alternating trials over nonalternating trials (main effect of Trial type ($F(1, 36) = 15.0$, $p < .001$, $\eta_p^2 = .29$); nonalternating trials: $M = 10.0$ sec, $SD = 4.55$; alternating trials: $M = 11.6$ sec, $SD = 4.42$). This preference becomes more pronounced in the final block (interaction between Trial type and

Figure 3 Schematic illustration of test trials. (a) Example of a nonalternating test trial (here, direct repetitions of token 3). (b) Example of an alternating trial (alternations between tokens 1 and 8, starting here with token 1).

Block: ($F(1, 36) = 4.27, p = .046, \eta_p^2 = .11$). Crucially, there is a 3-way interaction between Trial type, Block, and Group ($F(1, 36) = 4.37, p = .044, \eta_p^2 = .11$). As Figure 4 illustrates, when we split the data by Group, the unimodal group only reveals a preference for alternating trials in the second block (i.e., a significant interaction between Block and Trial type: $F(1, 18) = 8.99, p = .008, \eta_p^2 = .33$), while the bimodal group preferred alternating trials throughout test (i.e., significant effect of trial type: $F(1, 18) = 5.12, p = .036, \eta_p^2 = .22$). Subsequent paired *t*-tests for the unimodal group confirm that in the first block they showed no preference ($t(18) = 0.088, p = .93$), while in the second block, they preferred alternating trials ($t(18) = 5.06, p < .001$).

DISCUSSION AND CONCLUSION

Our results show that the initial frequency distribution of objects can affect infant category formation. This cannot be attributed to differences in attention during familiarization. Yet, it was only in the first test block that we observe the hypothesized group differences. Whereas the bimodal group continued to prefer alternating trials, the unimodal group changed from showing no preference to (also) preferring alternating trials. It appears that with additional input, infant category formation can change accordingly such that at the end of the testing phase, infants from the unimodal condition are also discriminating between tokens. How can we explain this shift in preferences?

One possible explanation is that after a unimodal familiarization, infants initially construct one broad category that with additional exposures becomes further divided into two (sub)categories. After all, the test phase itself resembles more a bimodal distribution, as here the midpoints were not presented anymore. Moreover, it is only in the familiarization phase that we presented filler tokens, which point to the existence of other categories. Perhaps additional exposures to only some (that is, distant) tokens

Figure 4 Mean looking times in seconds, for nonalternating (solid gray) vs. alternating (checked) test trials. Data are split by familiarization condition and by block. Error bars reflect 1 *SE* from the mean.

from a visual continuum, coupled with a different context in which filler tokens are now absent (Barsalou, 1982), led infants to remain attentive to those trials that high-light token differences (e.g., alternating trials), while losing interest in repetitive trials. Such an explanation fits with findings showing that sensitivity to relevant feature dimensions can shift through additional exposure with different contexts (e.g., Quinn, Schyns, & Goldstone, 2006; Schyns et al., 1998).

Another possibility is that additional exposure leads to a change in object recognition; for instance, infants start to recognize them as individuals. Given that infants generally prefer visual sequences that are optimized for learning, that is, “neither too simple nor too complex” (Kidd, Piantadosi, & Aslin, 2012), it is possible that infants viewed the nonalternating sequences as too simple, and thus less captivating than sequences of different tokens (regardless of whether or not these tokens belong to different categories). Our findings do not allow us to disentangle both possibilities. What remains clear is that a bimodal distribution enlarges the differences between tokens from opposite sides of the continuum, while this is not immediately apparent with a unimodal distribution.

In this paper, we set out to test whether distributional learning is a plausible domain-general mechanism for categorization. Our results on object perception corroborate that this is likely so as it couples with positive evidence in the domains of speech perception (e.g., Maye et al., 2002) and face perception (Altvater-Mackensen et al., 2017). Yet although category formation rapidly emerges based on initial frequency distributions of observed objects, it is far from fixed as it remains susceptible to additional exposures. Clearly, categorization is an unfolding process and is continually shaped and updated through new experiences.

ACKNOWLEDGMENTS

This research was funded by the research priority program Amsterdam Brain & Cognition of the University of Amsterdam and a personal VENI grant (016.154.051) from the Dutch Organisation for Scientific Research (NWO), awarded to CJ. The authors would like to thank parents and infants for their cooperation. Part of this research was earlier presented as a poster at the XIX Biannual International Conference on Infant Studies, July 2014, Berlin, Germany. The authors confirm that there are no conflicts of interests.

REFERENCES

- Altvater-Mackensen, N., Jessen, S., & Grossmann, T. (2017). Brain responses reveal that infants' face discrimination is guided by statistical learning from distributional information. *Developmental Science*, 20, e12393.
- Aslin, R. N., & Newport, E. L. (2012). Statistical learning: From acquiring specific items to forming general rules. *Current Directions in Psychological Science*, 21, 170–176.
- Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, 10, 82–93.
- Best, C., & Jones, C. (1998). Stimulus-alternation preference procedure to test infant speech discrimination. *Infant Behavior and Development*, 21, 295.
- Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition*, 121, 127–132.

- Feldman, N. H., Myers, E. B., White, K. S., Griffiths, T. L., & Morgan, J. L. (2013). Word-level information influences phonetic learning in adults and infants. *Cognition*, 127, 427–438.
- 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.
- Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1, 69–78.
- Goudbeek, M., & Swingle, D. (2006). Saliency effects in distributional learning. In P. Warren & C. I. Watson (Eds.), *Proceedings of the 11th Australian international conference on speech science & technology*, pp. 478–482.
- Hessels, R. S., Niehorster, D. C., Kemner, C., & Hooge, I. T. (2017). Noise-robust fixation detection in eye movement data: Identification by two-means clustering (I2MC). *Behavior Research Methods*, 49, 1802–1823.
- Johnson, M. H., Senju, A., & Tomalski, P. (2015). The two-process theory of face processing: Modifications based on two decades of data from infants and adults. *Neuroscience & Biobehavioral Reviews*, 50, 169–179.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS One*, 7, e36399.
- 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.
- Mandler, J. M. (2000). Perceptual and conceptual processes in infancy. *Journal of Cognition and Development*, 1, 3–36.
- Maye, J., Weiss, D. J., & Aslin, R. N. (2008). Statistical phonetic learning in infants: Facilitation and feature generalization. *Developmental Science*, 11, 122–134.
- Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition*, 82, B101–B111.
- Medin, D. L., & Coley, J. D. (1998). Concepts and categorization. In J. Hochberg, & J. E. Cutting (Eds.), *Perception and cognition at century's end: Handbook of perception and cognition* (pp. 403–439). San Diego, CA: Academic Press.
- Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, 32, 89–115.
- Nosofsky, R. M. (1988). Similarity, frequency, and category representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 54–65.
- Oakes, L. M., & Ribar, R. J. (2005). A comparison of infants' categorization in paired and successive presentation familiarization tasks. *Infancy*, 7, 85–98.
- Oakes, L. M., & Spalding, T. L. (1997). The role of exemplar distribution in infants' differentiation of categories. *Infant Behavior and Development*, 20, 457–475.
- Plunkett, K., Hu, J., & Cohen, L. B. (2008). Labels can override perceptual categories in early infancy. *Cognition*, 106, 665–681.
- Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. L. (1993). Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception*, 22, 463–475.
- Quinn, P. C., Schyns, P. G., & Goldstone, R. L. (2006). The interplay between perceptual organization and categorization in the representation of complex visual patterns by young infants. *Journal of Experimental Child Psychology*, 95, 117–127.
- Rosch, E. (1978). Principles of categorization. In E. Rosch, & B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Schyns, P. G., Goldstone, R. L., & Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, 21, 1–17.
- Shepard, R. N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function I. *Psychometrika*, 27, 125–140.
- Sloutsky, V. M. (2003). The role of similarity in the development of categorization. *Trends in Cognitive Sciences*, 7, 246–251.
- Smith, L. B., & Heise, D. (1992). Perceptual similarity and conceptual structure. In B. Burns (Ed.), *Percepts, concepts and categories: The representation and processing of information* (pp. 233–272). New York, NY: Elsevier.
- ter Schure, S. M. M., Junge, C. M. M., & Boersma, P. P. G. (2016). Semantics guide infants' vowel learning: Computational and experimental evidence. *Infant Behavior and Development*, 43, 44–57.
- Waxman, S. R., & Gelman, S. A. (2009). Early word-learning entails reference, not merely associations. *Trends in Cognitive Sciences*, 13, 258–263.

- Waxman, S. R., & Lidz, J. L. (2006). Early word learning. In D. Kuhn, & R. S. Siegler (Eds.), *Handbook of child psychology: Vol. 2, Cognition, perception, and language* (pp. 299–335), Hoboken, NJ: Wiley.
- Yeung, H. H., & Nazzi, T. (2014). Object labeling influences infant phonetic learning and generalization. *Cognition*, 132, 151–163.
- Yoshida, K. A., Pons, F., Maye, J., & Werker, J. F. (2010). Distributional phonetic learning at 10 months of age. *Infancy*, 15, 420–433.
- Younger, B. A. (1985). The segregation of items into categories by ten-month-old infants. *Child Development*, 56, 1574–1583.
- Younger, B. A., & Cohen, L. B. (1983). Infant perception of correlations among attributes. *Child Development*, 54, 858–867.