



PAPER

Information from multiple modalities helps 5-month-olds learn abstract rules

Michael C. Frank,¹ Jonathan A. Slemmer,² Gary F. Marcus³ and Scott P. Johnson⁴

1. Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, USA

2. Department of Psychology, Ohio State University, USA

3. Department of Psychology, New York University, USA

4. Department of Psychology, University of California, Los Angeles, USA

Abstract

By 7 months of age, infants are able to learn rules based on the abstract relationships between stimuli (Marcus et al., 1999), but they are better able to do so when exposed to speech than to some other classes of stimuli. In the current experiments we ask whether multimodal stimulus information will aid younger infants in identifying abstract rules. We habituated 5-month-olds to simple abstract patterns (ABA or ABB) instantiated in coordinated looming visual shapes and speech sounds (Experiment 1), shapes alone (Experiment 2), and speech sounds accompanied by uninformative but coordinated shapes (Experiment 3). Infants showed evidence of rule learning only in the presence of the informative multimodal cues. We hypothesize that the additional evidence present in these multimodal displays was responsible for the success of younger infants in learning rules, congruent with both a Bayesian account and with the Intersensory Redundancy Hypothesis.

Introduction

The ability to learn abstract regularities from a limited set of particulars is a powerful cognitive tool that comes into play in tasks as disparate as recognizing objects (Ullman, 1996), reasoning under uncertainty (Tenenbaum & Griffiths, 2001), and learning the rules of language (Pinker, 1989). Traditional empiricist views have long associated the time course of development with the stages of induction, assigning the gathering of information to infancy and the creation of more abstract knowledge to later childhood (Locke, 1964/1690; Piaget, 1952); however, more recent research has revealed evidence of abstraction away from perceptual particulars even in young infants (Kuhl, Williams, Lacerda, Stevens & Lindblom, 1992; Marcus, Vijayan, Bandi Rao & Vishton, 1999; Saffran, Pollak, Seibel & Shkolnik, 2007). A major goal of recent developmental research has been determining the sources of this abstraction, through investigations of the phylogenetic origins of these abilities (e.g. Ramus, Hauser, Miller, Morris & Mehler, 2000) and through research into the learning mechanisms available to young infants (e.g. Saffran, Aslin & Newport, 1996).

An explicit demonstration of infants' early abstraction abilities was given by Marcus *et al.* (1999). They familiarized 7-month-old infants to 2 minutes of syllable

strings like *ga ti ga* or *li na li*, where each string was of the form ABA (the first and last syllables were the same). At test, infants listened longer to strings instantiating a novel rule (ABB) than those instantiating the familiar rule, even though all strings were composed of syllables, such as *wo fe wo*, that had not been used during the training phase of the experiment. Marcus *et al.* concluded that the infants had successfully learned an abstract rule that was unbound to the particulars of the training stimuli. As with other artificial language learning paradigms used with both adults and children (Mintz, 2002; Saffran, Aslin & Newport, 1996; Saffran, Newport & Aslin, 1996; Smith, 1966), the rule learning paradigm serves as an excellent test case for further investigation of the mechanisms underlying infants' success in this type of task.

More recent research using this paradigm has focused on the perceptual and cognitive domains in which this type of mechanism can operate. Interest in the question of domain-specificity has largely been driven by the possibility of rule learning as a possible mechanism for language learning (Peña, Bonatti, Nespor & Mehler, 2002). However, even if the learning mechanisms responsible for successes in the Marcus *et al.* (1999) study are domain general, characterizing how they operate will still be of interest in understanding the inductive tools available to young infants.

Address for correspondence: Michael C. Frank, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139, USA; e-mail mcf Frank@mit.edu

In fact, the picture that has emerged from a variety of recent studies is that infants are able to induce abstract rules from a variety of materials across domains and modalities, but that the difficulty of rule extraction varies widely across different stimuli. For example, Saffran *et al.* (2007) reported that infants were successful in learning abstract rules in a visual stimulus set consisting of pictures of dogs presented simultaneously. Johnson, Fernandes, Frank, Kirkham, Marcus, Rabagliati and Slemmer (2008) reported a more mixed series of results using looming visual shapes presented sequentially, as with the auditory materials in the original Marcus *et al.* (1999) studies: 8-month-olds were able to learn and discriminate ABB rules from ABA but not AAB rules, while 11-month-olds were able to discriminate ABB from AAB as well. In contrast to the Marcus *et al.* studies, neither age group succeeded when they were trained on strings of the form ABA. These results suggest that acquiring rules instantiated in abstract visual shapes is more difficult than acquiring rules instantiated in speech, and that not all three-item identity rules are equally easy to learn. Finally, Marcus, Fernandes and Johnson (2007) showed that infants were not able to learn ABA or ABB rules when they were instantiated in (inherently sequential) auditory materials such as tones, timbres, and animal sounds, but succeeded when they were trained using speech rules and tested on materials in each of these domains.

In the context of these studies, we ask a separate but related question: can rule extraction be facilitated by training that uses multimodal stimuli? This issue is of interest for a number of reasons. First, multimodal information has been shown to be useful in a wide variety of perceptual and associative learning tasks (Bahrack, Flom & Lickliter, 2002; Bahrack & Lickliter, 2000; Flom & Bahrack, 2007; Gogate & Bahrack, 1998; Kirkham, Slemmer, Richardson & Johnson, 2007), but it is not known whether multimodal information supports more complex tasks such as rule learning. Second, the presence of multiple, redundant cues often allows infants to succeed in particular learning tasks earlier than they might in the presence of unimodal information (Flom & Bahrack, 2007; Gogate & Bahrack, 1998; Kirkham *et al.*, 2007). Thus using multimodal stimuli might allow younger infants to succeed in this task, suggesting that the mechanisms responsible for rule learning are present earlier than had previously been documented. Finally, the role of multimodal information in facilitating rule learning may be of interest in determining the nature of the underlying mechanisms.

In the following studies we investigate these questions. We tested whether 5-month-olds are able to learn ABB and ABA rules from input consisting of coordinated strings of colored, looming shapes and speech syllables (Experiment 1). Then, in two control experiments, we tested whether input in either modality is uniquely responsible for infants' success in learning the rules in

Experiment 1. In Experiment 2, we tested visual stimuli separated from speech, and in Experiment 3, we examined speech stimuli coordinated with an uninformative shape cue.

General methods

Participants

Ninety-six 5-month-olds composed the final sample ($N = 32$ for each experiment). Eight additional infants were observed but excluded from the analyses due to fussiness (four) or persistent inattention to the stimuli (four). We recruited infants by letter and telephone call from hospital records, commercial databases, and birth announcements; all infants were full term. Parents received a small gift as a token of their participation.

Apparatus and stimuli

Stimuli were presented using a Macintosh computer running Macromedia Director and a 53 cm color monitor. An experimenter viewed the infant over a closed-circuit television camera and coded looking times online by pressing a key when the infant was looking. The experimenter was blind to the stimulus being presented on screen. The experimenter and parent wore headphones and listened to music to avoid hearing the speech stimuli in Experiments 1 and 3.

Visual shape stimuli in Experiments 1 and 2 were identical to those in Johnson *et al.* (2008) and consisted of 12 geometric shapes (gray octagon, red square, green chevron, cyan diamond, blue bowtie, magenta 4-pointed star, orange triangle, yellow circle, white 5-pointed star, turquoise cross, pink clover, purple crescent). These shapes were presented one by one on a black background as in Kirkham, Slemmer and Johnson (2002). Each shape increased in size from 4 to 24 cm in height (2.4–14.6° visual angle) over the course of 1 s. Auditory stimuli in Experiments 1 and 3 were recordings of the syllables *ba*, *de*, *di*, *ga*, *je*, *ji*, *ko*, *le*, *li*, *po*, *we*, and *wi* (mean duration: 338 ms), identical to those used by Marcus *et al.* (1999). Each sequence lasted 3 s, followed by 1 s of background (silent).

All stimuli were organized into sequences of the form ABB or ABA, with six stimuli assigned to create three training sequences (e.g. *ba de de*, *di ga ga*, and *je ji ji*) and six stimuli assigned to create three novel test sequences which had not been heard during training (e.g. *ko le le*, *li po po*, and *we wi wi*). Sequences were randomly ordered during habituation and test with the constraint that sequences would not repeat immediately.

Procedure

Infants sat on a parent's lap during the studies, approximately 95 cm from the stimulus presentation monitor.

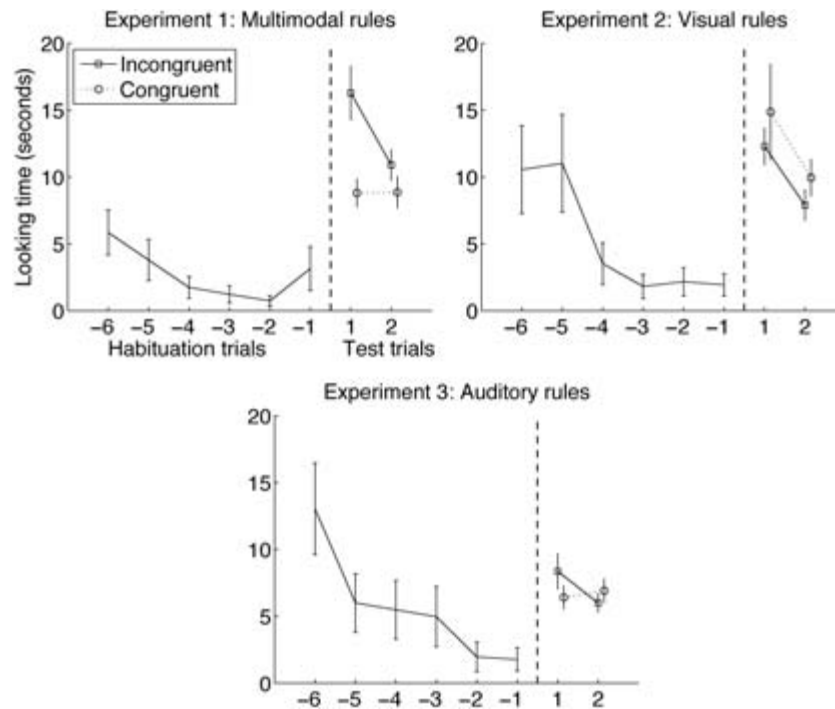


Figure 1 Results from Experiments 1–3. Graphs plot mean looking time (error bars indicate standard error of the mean) versus trial type. On the left side, the last six trials before habituation are shown; to the right of the dotted line, looking times to rule-congruent and rule-incongruent test stimuli are shown (X values are jittered slightly to avoid overlap of error bars).

Infants were habituated to either an ABB or an ABA pattern (counterbalanced across infants). Each trial began with the presentation of an engaging attention-getter (an expanding and contracting ball that beeped in conjunction with its motion). Once the experimenter detected that the infant was looking at the monitor, the experimenter pressed a key and stimulus presentation began. In each trial, sequences of three items were presented in random order until the trial ended. Sequences within each trial were chosen uniformly from the three training sequences described above without immediate repetitions. When infants turned away from the monitor, the experimenter released the key press and stimulus presentation was immediately paused. If the infant returned attention toward the screen the experimenter again pressed the key and stimulus sequences resumed (only looking time, not including time spent looking away, was included in the measured length of a habituation trial). Trials were terminated after 2 s of continuous looking away or a maximum of 90 seconds.

Habituation continued until 12 trials elapsed or looking times across four trials declined to less than 50% of looking time during the first four trials. After the habituation phase was ended, infants saw four test trials alternating between the pattern they had seen during training and the opposite pattern; both patterns were instantiated in entirely novel stimuli. Test trial order was counterbalanced across infants (half saw a novel pattern first while half saw a familiar pattern first).

Experiment 1: Multimodal rules

In the first experiment, 5-month-old infants ($N = 32$, 15 girls and 17 boys, M age = 155.1 days) were exposed to a multimodal pattern, in which shapes and syllables were presented simultaneously with both reflecting the same underlying rule (e.g. for an ABA rule, the pattern would be *ba*-octagon, *de*-square, *ba*-octagon).

Results and discussion

Looking times to novel and familiar patterns at test (all stimuli were novel) for this and the other experiments are shown in Figure 1. At test, infants looked significantly longer at the novel pattern than the familiar one, according to a paired, two-tailed t -test ($t(31) = 3.54$, $p = .001$), suggesting that they had learned some abstract rule during the training portion of the study and thus dishabituated when presented with a novel rule during test.¹ There was no significant difference in novelty preference between infants trained on ABA or ABB stimuli in a two-sample, unpaired t -test ($M = 63.1\%$ and 57.5% , respectively, $t(30) = 1.01$, ns), suggesting that infants learned the appropriate pattern in both conditions. Altogether, 24 of 32 infants

¹ The paired t -test is appropriate for looking time data (which may be log-normally, rather than normally, distributed) because it operates over the difference of the looking to incongruent and congruent stimuli, which is in fact normally distributed.

in this experiment showed a novelty preference; this figure differs significantly from chance in a two-tailed sign test ($p = .007$). To ensure that this effect was independent of whether infants reached habituation criterion, we tested whether the novelty preference of the 29 infants that met the habituation criterion was significant, and we found that it was ($t(28) = 4.06, p < .001$).

Experiment 2: Visual rules

In the next experiment, we tested whether 5-month-old infants ($N = 32$, 16 girls and 16 boys, M age = 154.7 days) were able to learn the same type of abstract rule from purely visual input as they are able to in some cases at 8 months (Johnson *et al.*, 2008). Infants were habituated to visual stimuli in which rules were instantiated only in shapes with no auditory input (e.g. for an ABA rule, the pattern would be octagon, square, octagon). All methods were otherwise identical to those for Experiment 1.

Results and discussion

At test, infants showed no significant preference for either the novel or familiar pattern in a paired t -test ($t(31) = -1.10, ns$), suggesting that visual information alone was not sufficient for these infants to learn the abstract pattern instantiated in the training stimuli. Preferences for infants habituated to ABB and ABA did not differ significantly from one another in an unpaired t -test ($t(30) = 0.77, p = .45$). In this experiment, 18 of 32 infants showed a novelty preference; this figure did not significantly differ from chance in a two-tailed sign test ($p = .60$). Only 13 of the 26 infants that met the habituation criterion in this experiment showed a novelty preference, and their numerical preference was for the familiar rule, though not significantly so ($t(25) = -1.46, p = .16$).

Experiment 3: Auditory rules

In our final experiment, we tested whether 5-month-olds could learn rules from multimodal input where only one modality (in this case, the speech modality) was informative about the underlying rule; if infants in this condition were able to succeed in learning the rule, that result would suggest that the multimodal input in Experiment 1 was primarily useful as an attentional cue, rather than providing stronger evidence of the presence of a particular rule (we discuss these two possibilities further in the General discussion). We constructed a stimulus set that was identical to that of Experiment 1, except that there was only one shape (a gray octagon) which loomed in synchrony with the presentation of the patterned syllables (so that a sample ABA familiarization sequence would be *ba*-octagon *de*-octagon *ba*-octagon). Our population was a group of 5-month-old infants ($N = 32$, 18 girls and 14 boys, M age = 154.7 days).

Results and discussion

We found no significant differences in looking time between the familiar and novel stimuli in a two-tailed, paired t -test ($t(31) = 0.60, ns$), suggesting that the presence of uninformative looming shapes coordinated with patterned speech stimuli did not facilitate rule learning. In this experiment, 16 of 32 infants showed a novelty preference; this figure did not significantly differ from chance in a two-tailed sign test ($p = 1.00$). Repeating this test with only the 28 infants that met our habituation criterion similarly did not result in a significant novelty preference ($t(27) = 0.52, p = .61$).²

We next attempted a further test of the attentional account of our findings. Perhaps the positive evidence for learning that we observed in Experiment 1, but not in Experiments 2 and 3, stemmed from differences in looking times during the habituation phase of the experiment: Longer looking times may have facilitated learning. If this were so, perhaps the variability in both modalities caused infants to attend longer and hence allowed them to learn the rules. However, habituation times and number of trials to habituation among the three experiments did not differ significantly from one another (mean total habituation times [number of trials] for Experiments 1, 2, and 3 were 142.9 s [6.7], 149.7 s [7.6], and 132.0 s [7.7], respectively, $F[2, 93] = 0.76, p = .47$ for times and $F[2, 93] = 1.76, p = .18$ for number of trials), suggesting that this explanation was not supported by the data. Finally, a direct comparison of habituation times in Experiments 1 and 3 also failed to reach significance ($t(62) = 0.82, ns$), suggesting that the different outcomes in these experiments were not due to the failure of the stimuli to capture infants' attention.

Given that two of our three experiments produced no significant novelty preference, we next asked whether the results of these experiments differed significantly from one another, or whether the pattern of performance we observed was simply an artifact of testing for significance individually in each experiment. To test this question, we conducted an analysis of variance in novelty preference (looking to novel test items / total looking at test) with experiment as our single factor. This analysis showed a significant effect of experiment on novelty preference ($F[2, 93] = 3.94, p = .02$), suggesting that the results of Experiment 1 were reliably different from the results of the other two experiments.

² In order to ensure that the gray octagon did not distract infants from the underlying regularity in Experiment 3, we conducted an additional control experiment. We used the same auditory stimuli as Experiment 3 but employed a static checkerboard as the accompanying visual stimulus (e.g. Colombo & Horowitz, 1986; Cooper & Aslin, 1990). Our participants were a new group of infants ($N = 32$, 11 girls and 21 boys, M age = 153.5 days). During the test phase we found no significant difference in looking to the novel and the familiar pattern ($t(31) = 0.87, ns$), suggesting that as in Experiment 3, information from speech alone was not sufficient for the 5-month-olds to extract the abstract pattern.

General discussion

Here we have presented evidence that an informative, multimodal stimulus presentation allowed 5-month-olds to extract abstract rules from short familiarization sequences. In two control experiments, we showed that neither auditory nor visual information alone were sufficient to allow 5-month-olds to learn the same regularity, even when a synchronized visual cue provided an uninformative but engaging multimodal cue. These results suggest two conclusions. First, they confirm and extend earlier work suggesting that early abstraction is facilitated by the presence of multiple, redundant information sources (Bahrick *et al.*, 2002; Bahrick & Lickliter, 2000; Flom & Bahrick, 2007). Second, in combination with the results of Johnson *et al.* (2008) and Saffran *et al.* (2007), our results provide further confirmation that abstract rule learning operates with different degrees of difficulty (and at different times in development) over stimuli in several domains and modalities, though speech may play some special role in rule learning that is as yet unspecified (Marcus *et al.*, 2007).

As suggested by the results of Dawson and Gerken (2006, 2008) indicating that younger infants may be able to learn musical rules that older ones cannot, mechanisms of rule learning may even become tuned to the regularities of particular modalities. On this account, 7-month-olds may learn rules with only speech (rather than multimodal) input because they have accumulated more evidence than 5-month-olds that speech tends to be organized in a rule-based fashion. In much the same way that infants lose sensitivity to non-native phonetic contrasts (Werker & Tees, 1984) and musical rhythms (Hannon & Trehub, 2005) while gaining sensitivity to those relevant to their own language and culture, infants' ability to detect rules in domains that are not rule-governed may decrease as their sensitivity increases in the domain of language. Additional research into the developmental trajectory of rule learning across domains is needed to evaluate this possibility.

How does the multimodal presentation employed in Experiment 1 facilitate rule learning in younger infants? On a Bayesian interpretation, intermodal cues could facilitate learning by providing learners with more evidence in favor of a particular rule. For instance, the model of Tenenbaum and Griffiths (2001) uses the *size principle* to restrict generalization. The size principle states that the probability of a data point under a hypothesis is inversely proportional to how general that hypothesis is – the more specific the hypothesis, the more likely a learner is to observe any data point consistent with it. In this kind of framework, a hypothesis about structure in two modalities is more specific than a hypothesis that covers only one, thus rational learners should get more evidence by observing an example of a multimodal rule than they would by observing an example of a unimodal rule.

For instance, the repetition in a string like *wo fe fe* could conceivably have occurred by chance (e.g. if strings

are made from 8 syllables chosen randomly, there is a 1 in 8 chance that any syllable is followed by itself). In contrast, the dual repetition in the string *wo*-octagon *fe*-square *fe*-square is much less likely (if there are 8 shapes as well, there is only a 1 in 64 chance that a shape/syllable combination is followed by itself). Thus, if the task of the learner is to evaluate how well hypotheses about the world are supported by available evidence, redundant information from two modalities gives far more evidence for the repetition hypothesis than does information from a single modality. On this kind of account, younger infants – such as the 5-month-olds tested here – may simply require more evidence (a 'more suspicious coincidence') than older infants for learning to occur. This account also opens the door for explanations of the effects of increased stimulus variability on learning (Gomez, 2002): greater variability in a stimulus creates a larger space of possible outcomes, thus rendering any particular outcome less probable. Put another way, if we had simply increased the number of shapes in Experiment 2 to 64 (thus making the chance of a repetition 1 in 64), perhaps infants would have succeeded in learning as well as they did in Experiment 1. A key test of the particular Bayesian perspective expressed here will be investigating whether increasing evidence (via variability or other methods) will have the same effects on learning as multimodal presentation does.

While this Bayesian account represents a new theoretical viewpoint that makes a variety of predictions regarding rule-learning tasks (Frank, Ichinco & Tenenbaum, 2008; Gerken, 2006; Tenenbaum & Griffiths, 2001), we believe that both our data and the Bayesian perspective more generally are also consistent with the *Intersensory Redundancy Hypothesis*, or IRH (Bahrick *et al.*, 2002; Bahrick & Lickliter, 2000; Flom & Bahrick, 2007). Under this account, the multimodal information available in Experiment 1, but not in Experiments 2 and 3, increased the salience of the amodal regularity (the repetition) in both modalities, allowing them to learn it more effectively. Rather than providing conflicting accounts of the same phenomenon, the two theories – the Bayesian account and the IRH – may operate at different levels of description (Marr, 1982). The Bayesian account provides an explanation of why a rational statistical learner might see the stimuli in Experiment 1 as giving more evidence than those in Experiment 2 for the same rule, while the IRH gives an account of the mechanisms at work in individual infants.

However, the Bayesian account and the IRH differ in at least two ways. First, the IRH makes a different prediction in the increased-variability case mentioned above, predicting that multimodal stimuli may have a different effect on learning than that caused by increasing the amount of unimodal evidence present in the stimulus. Second, we know of no current exploration in the Bayesian concept learning framework of one of the primary phenomena addressed by the IRH: temporal synchrony (Bahrick & Lickliter, 2000; Flom & Bahrick, 2007).

Exploring the facilitatory effects of temporal synchrony on learning will require elaboration of the Bayesian framework to deal with coincidences that occur within a continuous dimension (time) as well as in samples from a discrete set of stimuli as we have discussed here. Thus, we hope that future work – experimental and computational – will further investigate both the points of congruence and the differences between these two theoretical frameworks.

Acknowledgements

This research was supported by NSF grant BCS-0418103 and NIH grants R01-HD40432 and R01-HD048733. We gratefully acknowledge the efforts of the infants and parents who participated in the studies. We also thank Myque Harris, Kristin Bellanca, and Lauren Clepper for help recruiting the infant participants and three anonymous reviewers for helpful comments and discussion. This work is dedicated to the memory of our friend and colleague Jon Slemmer.

References

- Bahrnick, L.E., Flom, R., & Lickliter, R. (2002). Intersensory redundancy facilitates discrimination of tempo in 3-month-old infants. *Developmental Psychobiology*, **41**, 352–363.
- Bahrnick, L.E., & Lickliter, R. (2000). Intersensory redundancy guides attentional selectivity and perceptual learning in infancy. *Developmental Psychology*, **36**, 190–201.
- Colombo, J., & Horowitz, F.D. (1986). Infants' attentional responses to frequency modulated sweeps. *Child Development*, **57**, 287–291.
- Cooper, R.P., & Aslin, R.N. (1990). Preference for infant-directed speech in the first month after birth. *Child Development*, **61**, 1584–1595.
- Dawson, C., & Gerken, L. (2006, November). Differential processing of language and music learned during development. Paper presented at the Boston University Conference on Language Development, Boston, MA.
- Dawson, C., & Gerken, L. (2008). Learning to learn differently: the emergence of domain-sensitive generalization in the second six months of life. Manuscript submitted for publication.
- Flom, R., & Bahrnick, L.E. (2007). The development of infant discrimination of affect in multimodal and unimodal stimulation: the role of intersensory redundancy. *Developmental Psychology*, **43**, 238.
- Frank, M.C., Ichinco, D., & Tenenbaum, J. (2008). Principles of generalization for learning sequential structure in language. *Proceedings of the 30th Annual Meeting of the Cognitive Science Society*.
- Gerken, L. (2006). Decisions, decisions: infant language learning when multiple generalizations are possible. *Cognition*, **98**, B67–B74.
- Gogate, L.J., & Bahrnick, L.E. (1998). Intersensory redundancy facilitates learning of arbitrary relations between vowel sounds and objects in seven-month-old infants. *Journal of Experimental Child Psychology*, **69**, 133–149.
- Gomez, R.L. (2002). Variability and detection of invariant structure. *Psychological Science*, **13**, 431–436.
- Hannon, E.E., & Trehub, S.E. (2005). Tuning in to musical rhythms: infants learn more readily than adults. *Proceedings of the National Academy of Sciences*, **102**, 12639–12643.
- Johnson, S.P., Fernandes, K.J., Frank, M.C., Kirkham, N.Z., Marcus, G.F., Rabagliati, H., & Slemmer, J.A. (2008). Abstract rule learning for visual sequences in 8- and 11-month-olds. *Infancy* (in press).
- 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 spatio-temporal sequence learning in infancy. *Child Development*, **78**, 1559–1571.
- Kuhl, P.K., Williams, K.A., Lacerda, F., Stevens, K.N., & Lindblom, B. (1992). Linguistic experience alters phonetic perception in infants by 6 months of age. *Science*, **255**, 606–608.
- Locke, J. (1964/1690). *An essay concerning human understanding*. Cleveland, OH: Meridian Books.
- Marcus, G.F., Fernandes, K.J., & Johnson, S.P. (2007). Infant rule learning facilitated by speech. *Psychological Science*, **18**, 387–391.
- Marcus, G.F., Vijayan, S., Bandi Rao, S., & Vishton, P.M. (1999). Rule learning by 7-month-old infants. *Science*, **283**, 77–80.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Henry Holt.
- Mintz, T.H. (2002). Category induction from distributional cues in an artificial language. *Memory and Cognition*, **30**, 678–686.
- Peña, M., Bonatti, L.L., Nespor, M., & Mehler, J. (2002). Signal-driven computations in speech processing. *Science*, **298**, 604–607.
- Piaget, J. (1952). *The origins of intelligence*. New York: Norton.
- Pinker, S. (1989). *Learnability and cognition: The acquisition of argument structure*. Cambridge, MA: MIT Press.
- Ramus, F., Hauser, M.D., Miller, C., Morris, D., & Mehler, J. (2000). Language discrimination by human newborns and by cotton-top tamarin monkeys. *Science*, **288**, 349–351.
- Saffran, J.R., Aslin, R.N., & Newport, E.L. (1996). Statistical learning by 8-month-old infants. *Science*, **274**, 1926–1928.
- Saffran, J.R., Newport, E.L., & Aslin, R.N. (1996). Word segmentation: the role of distributional cues. *Journal of Memory and Language*, **35**, 606–621.
- 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.
- Smith, K.H. (1966). Grammatical intrusions in the recall of structured letter pairs: mediated transfer or position learning? *Journal of Experimental Psychology*, **72**, 580–588.
- Tenenbaum, J.B., & Griffiths, T.L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, **24**, 629–640.
- Ullman, S. (1996). *High-level vision: Object recognition and visual cognition*. Cambridge, MA: MIT Press.
- Werker, J.F., & Tees, R.C. (1984). Cross-language speech perception: evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development*, **7**, 49–63.

Received: 12 September 2007

Accepted: 30 April 2008