

A bottom-up view of toddler word learning

Alfredo F. Pereira · Linda B. Smith · Chen Yu

Published online: 28 June 2013
© Psychonomic Society, Inc. 2013

Abstract A head camera was used to examine the visual correlates of object name learning by toddlers as they played with novel objects and as the parent spontaneously named those objects. The toddlers' learning of the object names was tested after play, and the visual properties of the head camera images during naming events associated with learned and unlearned object names were analyzed. Naming events associated with learning had a clear visual signature, one in which the visual information itself was clean and visual competition among objects was minimized. Moreover, for learned object names, the visual advantage of the named target over competitors was sustained, both before and after the heard name. The findings are discussed in terms of the visual and cognitive processes that may depend on clean sensory input for learning and also on the sensory–motor, cognitive, and social processes that may create these optimal visual moments for learning.

Keywords Visual attention · Language comprehension · Word recognition

Children learn their first object names by linking a heard word to a seen thing. Contemporary theories all assume that the learning environment is noisy, with scenes containing several potential referents for a heard name. Different theories posit different mechanisms through which young learners reduce this uncertainty, including social cues to speaker intent (Baldwin, 1995; Tomasello & Akhtar, 1995), innate linking functions between linguistic categories and meanings (Booth & Waxman, 2009; Lidz, Waxman, & Freedman, 2003), and statistical mechanisms that aggregate word–object co-occurrences across multiple naming events (Frank, Goodman, & Tenenbaum, 2009; Smith & Yu, 2008; Xu & Tenenbaum, 2007). Here, we present new evidence on the nature of the learning environment at the sensory level, in terms of the

moment-to-moment visual information available to the learner about potential referents for a heard name. The findings raise questions about the starting assumption of rampant ambiguity in the early object-name-learning environment and suggest new hypotheses about how visual clutter and competition may limit early word learning.

Our interest in and approach to studying the dynamic visual correlates of object-name learning stem from four considerations. First, the everyday visual world not only offers many potential referents, but also is dynamically complex; objects in the scene move and change in relation to each other and in relation to the sensors as the perceiver also acts and moves. Second, a large literature studying toddler attention shows how this everyday context of a moving body and moving objects is attentionally challenging (e.g., Kanass, Oakes, & Shaddy, 2006). Indeed, sustained attention during play with multiple objects is used to assess individual differences in attentional functioning in typically and atypically developing toddlers (e.g., Lawson & Ruff, 2004). Third, a growing literature on atypical development indicates the comorbidity of sensory–motor, attention, and language delays (e.g., Iverson, 2010). These links are not well understood mechanistically. However, the significant changes in motor behavior that characterize the second year of life (e.g., Adolph & Berger, 2006) bring with them bodily instabilities and, as a result, large head and trunk movements (Berthenthal & von Hofsten, 1998). These movements directly affect the visual input and potentially destabilize attention and may create special challenges to object name learning. Finally, several recent studies have used head cameras to capture the moment-to-moment visual dynamics as toddlers engage in various activities (Aslin, 2009; Cicchino, Aslin, & Rakison, 2010; Yoshida & Smith, 2008). These studies show that toddlers' head-centered views during active play are not at all like adult views in that they are highly dynamic, with individual objects coming into and going out of view on times scales of seconds and fractions of seconds (Smith, Yu, & Pereira, 2011). All four considerations suggest the value of studying the ambiguity of naming movements from the perspective of the dynamic properties of visual experience.

A. F. Pereira (✉)
Universidade do Minho, Braga, Portugal
e-mail: alfredo.pereira@algoritmi.uminho.pt

L. B. Smith · C. Yu
Indiana University Bloomington, Bloomington, USA

One particular result from the prior head camera studies motivates our specific experimental question. Amidst the highly dynamic views that were found to characterize active toddlers' visual experiences were occasional less dynamic periods when, despite many objects being in near physical proximity to the child, there was just one object stably dominating the head camera image, being much larger in visual size because it was closer and unoccluded (Smith et al., 2011; Yoshida & Smith, 2008; Yu, Smith, Shen, Pereira, & Smith, 2009; 2009). We ask, Are these periods of stable, clean, nearly one-object views optimal sensory moments for the early learning of object names? To answer this question, toddlers' first-person views were recorded by a head camera as they played with several novel objects with a parent and as the parent spontaneously named those objects. The toddlers' learning of the object names was tested after play, and the visual properties of the head camera images during naming events associated with learned and unlearned object names were analyzed. On the basis of the prior head camera studies, the main dependent measures were the temporal profile of the named object's image size before, during, and after a naming and the same temporal profile for the unnamed competitors. We also measured the centering of the objects in the image, providing a dynamic profile of the spatial direction of attention with respect to the named target and unnamed competitors around moments of naming.

Method

Participants

Twelve toddlers (7 male, 16–25 months, $M=20$ months) were recruited to participate. Three additional children did not contribute data to the final analyses, either because of failure to tolerate the head camera or because of calibration difficulties.

Stimuli

Six novel objects (on average, about $9.5 \times 6.5 \times 5$ cm) were custom made from hardened clay to have unique shapes and textures. Each object was randomly paired with one name (*zeebee*, *tema*, *dodi*, *habble*, *wawa*, and *mapoo*), and the objects were organized in two sets of three. Within each set, one object was painted blue, one red, and one green.

Head camera

The mini head camera (KT&C model VSN500NH, f2.45, 768×494 pixels CCD resolution) was embedded in a custom headband and recorded a broad 97° visual field in the horizontal—approximately half of the visual field of infants

(Mayer & Fulton, 1993)—and 87° in the vertical. A prior calibration study (Yoshida & Smith, 2008) independently measured eye gaze direction and head direction during toy play and found that noncorrespondence between head and eye was generally infrequent (less than 17 % of frames) and brief (less than 500 ms; see also Smith et al., 2011; Yu et al., 2009). To place the head camera on the infant, one experimenter distracted the child, while the second placed the head camera on the head. The child was then directed to push a button on a pop-up toy, and the camera was adjusted such that the button at the moment it was pushed was centered in the head camera image. Additional third-person cameras were used to record the play session and to record the experimenter and the child during testing.

Experimental room

Parents and infants sat across from each other at a small table ($61 \times 91 \times 64$ cm) that was illuminated from above. The average distance of the infant's eye to the center of the table was 43.2 cm. Parent and toddler wore white clothing, and the walls, table, and floor were also white so that shadows were minimized.

Procedure

Prior to the play task, the parents were instructed as to the names of the six novel objects and were asked to use these names during play. To remind the parents of the object labels, the labels and object pictures were attached to the boxes from which parents retrieved the two toy sets. Parents were *not* told that their task was to teach the names or that infants would be later tested. They were only told to encourage their child's interaction with the objects in as natural a manner as possible. Parents were alone in the room with their child during the play period. There were four toy play trials, two with each set of three objects, lasting 1.5 min each. The start and stop of each play period was cued by an auditory signal. The parent's voice was recorded using a noise reduction microphone.

After the play trials, an experimenter entered the room and tested the toddlers' knowledge of the object names. On each test trial, three objects were placed on a tray, 44 cm wide, such that one object was to the extreme right, one to the extreme left, and one at midline. The experimenter held the tray away from the infant, looked continually into the infant's eyes, never at the objects (as confirmed by video recording), and said "Show me the ___! Get the ___!" and then moved the tray forward for the infant to select an object. Each of the six object names was tested twice (with all three objects tested once before any object was tested a second time). The distractors on each trial were randomly selected from the other play objects, with the following

constraints: All objects served as distractors equally often; each trial was composed of one red, one blue, and one green object; and the distractors used for any target differed on the two testing trials. The location of the correct object varied (via a Latin square) across trials for each infant.

Coding

A naming event was defined as any whole parent utterance (e.g., “*What are you doing with that habble?*”) that contained an object name. A silence duration of more than 0.4 s was used to mark the temporal boundaries of utterances, and human coders then identified utterances that included the object names. Agreement for two coders for a randomly selected set of utterances exceeded 90 %, and all disagreements were resolved by the two coders relistening to the audio recordings.

The head camera video was sampled at 10 Hz, and head camera images were analyzed frame-by-frame for the 10 s prior, during, and for 10 s after each naming event, yielding approximately 640 data points (frames) for each naming event. Measures of the visual properties were taken for each of the three play objects in the defined window using a custom image-analysis software (see Yu et al., 2009): (1) the image size of each of the three objects, measured by proportion of object pixels in the image, and (2) the centering of each object in the image, measured by computing the average distance of all object pixels to the image center and expressing that average distance as a proportion of the head camera image’s half diagonal—that is, a fully centered object pixel corresponds to zero centering, and a head camera image corner pixel has a centering value of one. All objects were the same physical size; thus, image size and overlap with center varies with infant and object movements. For the statistical analyses, the 10-Hz time series were averaged within the utterance containing the naming event and within 1-s windows for each of the 10 s prior to and after the naming utterance.

The toddler’s performance at test was scored by a naïve human coder who did not know the correct choice and who made an all-or-none decision as to the selected object on each object name test trial. A second scorer scored a randomly selected 25 % of the test trials, and the level of agreement exceeded 94 %. An object name was defined as “learned” if the toddler correctly selected it on two of the two testing trials; otherwise, the object name was considered “not learned.”

Results

Parents produced each of the six object names, on average, 9.7 times ($SD=4.4$). At test, infant choices indicated that, on average, 1.58 names were learned (range across the 12 toddlers, 1–6); overall, this level of success exceeds that expected

by random choice (0.67 correct names), $t(11)=3.19$, $p<.01$ (two-tailed). However, the key issue is not whether infants could learn some object names but, rather, the visual properties of the individual naming events that supported this learning. Accordingly, naming events were partitioned into those associated with learned versus unlearned object names. This is a noisy partition, since not all naming events associated with learned object names may have contributed to learning. The mean number of naming utterances per parent associated with each learned object name was 11.0 ($SD=6.1$) and was more than the number of naming utterances associated with unlearned object names, 8.0 ($SD=4.2$), $t(70)=-2.27$, $p<.05$. The average duration of utterances containing a name was 1.25 s, $SD=0.61$ s, and was slightly less than the average duration of utterances associated with unlearned object names of 1.38 s ($SD=0.73$ s), $t(637)=-2.13$, $p<.05$. Both of these factors could contribute to learning; the key question for this study, however, concerns the dynamics of the visual properties of the naming events.

Do the visual properties of naming events associated with learned names differ from those associated with nonlearned names? To answer this question, the two dependent measures, the image size of the objects and their centering in the image, were analyzed for 10 s before and 10 s after each naming event with the critical questions concerning the temporal profiles of these properties for the named target and for the other objects, the potential competitors. The analyses examine the properties of the head camera image for a 20-s window around a naming utterance. More than one naming utterance (for the same or different objects) could potentially be contained in the same 20-s window around a single naming utterance, yielding overlapping 20-s windows for two different naming events. These were included in the analyses because they were relatively infrequent and did not differ for learned and unlearned object names. The proportion of naming events that overlapped each other within the 20-s window was 8.6 % for learned object names and 10.9 % for unlearned object names.

The analyses were conducted on a total of 639 naming events (209 and 430 for learned and unlearned object names, respectively) and used the methodology of growth curve analysis (GCA). Separate GCAs were conducted for image size and object centering. GCA is a type of hierarchical linear modeling concerned with capturing time effects under assumptions of a continuous stochastic process and is structured hierarchically at least two levels (see Mirman, Dixon, & Magnuson, 2008). At level 1, the growth curve for each dependent variable is modeled by a linear regression using time as a predictor. The regression model can include zero-order (intercept), first-order (slope), and higher-order polynomial *time* terms. Because the polynomial terms are naturally collinear, they were transformed into orthogonal polynomials so that the contribution of each polynomial term could be

assessed independently of the others. The level 2 model considers the level 1 model as potentially explainable by a linear regression of population averages, fixed effects (typically, the effects of interest), and random effects and thus serves the role in the analyses of the more typical analysis of variance. To build the level 1 and level 2 models, we followed the methodology of Baayen, Davidson, and Bates (2008). A model comparison approach based on a likelihood ratio test was used, and models were checked for possible overfitting by examining the residuals of any random effects and the correlations between fixed effects. Visual inspection of temporal profile for object size and centering measures revealed a clear U-curve, inverted for object size and U-shaped for centering, with the maximum (object size) or minimum (centering) point at the naming utterance. Consequentially, we explored, for level 1, models that included an intercept, a linear, and a quadratic time term. In order to account for individual and stimuli differences, we considered a *participant* random effect and a separate *object label* random effect (i.e., these were crossed random effects); we did not include interactions between time and participant or object label. The level 2 model was constructed in two steps: First, we used a series of model comparisons to determine different random effect structures (intercept, linear, and quadratic terms) for participant and object label effects; second, we added a full two-way interaction between the fixed effects of interest, (1) *named object* (target/competitor) and (2) *learning* (learned/unlearned). Model parameters were estimated using the lme4 package (Bates, 2005, 2012; available in R, R Development Core Team, 2008). Fixed effects were contrast-coded, and *p*-values for model parameter estimates computed using a Markov chain Monte Carlo (MCMC) simulation method (see Baayen, Davidson, & Bates, 2008).

In preview, the main conclusions that arise from the analyses, evident in Fig. 1, are the following. First, whether or not the name is learned, the visual properties of the named targets differed from those of unnamed ones; specifically, for both learned and unlearned object names, the named target had an image size advantage over competitors and was more centered in the visual field than were the unnamed competitor objects. Second, named targets that were learned differed from named targets that were not learned in the magnitude of the difference in these visual properties between the named target and the other, competitor, objects. Specifically, naming events for learned names showed a larger difference between named target and competitors, with the implication of less visual competition, than did the named targets that were not learned.

Object size

The temporal profiles for image size for the target and (the average of) the competitors are shown in Fig. 1c, d, and the main results of the GCA are given in Table 1. The GCA

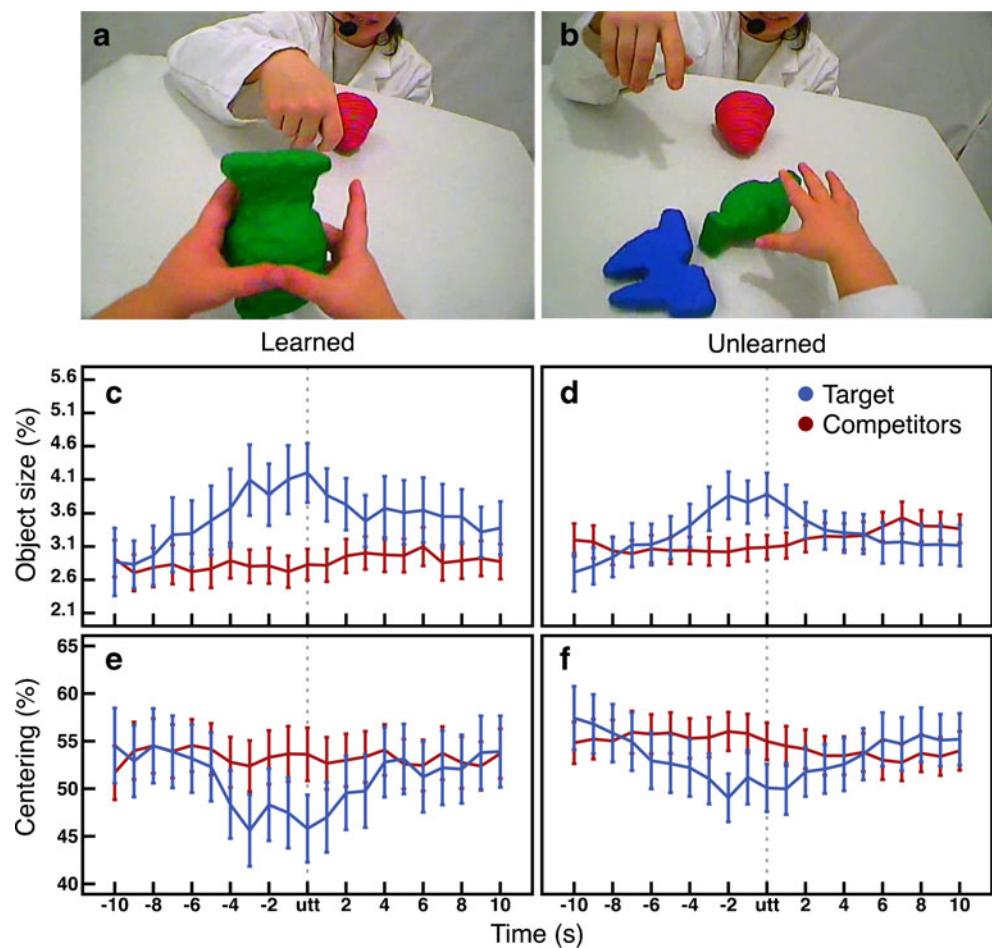
yielded a best-fit model with a quadratic, $B=-0.58$, $p<.001$, time term, indicating a rise and then fall of image size before and after a naming event and, thus, a clear dynamic link between image size and naming events. The GCA also yielded an average image size advantage for the named target versus the unnamed competitors, $B=0.51$, $p<.01$, but no main effect for learned versus unlearned words, $B=0.08$, $p<.093$. Critically, the analysis yielded a reliable learning \times named object interaction, $B=0.92$, $p<.001$, since the named target's image size advantage was greater for learned than for unlearned object names. The maximum correlation between fixed effects was moderate, $r=.41$. The analysis also revealed a random intercept per participant and a random intercept per object label. These indicate individual differences and stimulus differences (reflecting stimulus-specific differences in how the infants held and interacted with the objects). The main conclusion, as apparent in Fig. 1c, d, is that naming events associated with learned object names, more than those associated with unlearned object names, are characterized by temporal profile in which the image size for the named target is larger than that for the unnamed competitors.

To determine when, in the time series, the named target diverged in image size from the mean of the competitors, we determined the first and last significant difference in a series of ordered pairwise *t*-tests (Allopenna, Magnuson, & Tannenhaus, 1998). For naming events associated with learning, the target advantage was stable and enduring: Image size was reliably different for the target versus competitors at 6 s prior to the naming event and persisted until 5 s after the event. For naming events associated with unlearned object names, there was also a target advantage, but it was much briefer; image size was reliably different for the target versus competitors only at 3 s prior to the naming event and persisted until 1 s after the event. In sum, for naming events associated with learning, the named object was more visually dominant than the competitors—larger in the field because it was closer and unoccluded—and this dominance was sustained over time.

Centering

The temporal profiles for centering for target and (the average of) the competitors are shown in Fig. 1e, f, and the main results of the GCA are also given in Table 1. The GCA for this measure yielded a best-fit model with a linear, $B=-1.7$, $p<.05$, and a quadratic, $B=4.1$, $p<.001$, time term. Centering, like image size, rises up to the naming event and then falls after the naming event. There was a reliable effect of named object, with an advantage in centering for the named target over competitors, $B=-1.6$, $p<.001$, and also an effect of learning, $B=-1.8$, $p<.001$. Similar to the object size measure, a significant two-way interaction of learning \times named object indicates that the target advantage in centering over the unnamed competitors is larger for naming events associated with learning,

Fig. 1 Example of a visual scene, while the parent labeled a target referent (green), for a referent for which the child learned the object name (a) and for a target referent (blue) for which the child did not learn the object name (b). Temporal profiles for object size—measured as proportion of the image—for the target and (the average of) the competitors are shown for naming events associated with learned words (c) and for naming events associated with unlearned words (d). Temporal profiles for centering—measured as average object pixel distance to center expressed as proportion of half-diagonal—for target and (the average of) the competitors are shown for naming events associated with learned words (e) and for naming events associated with unlearned words (f)



$B = -2.7$, $p < .001$. The maximum correlation between fixed effects was moderate, $r = .44$. The analysis also yielded a random intercept per participant and a random intercept per object label, again showing individual differences and stimulus differences in centering. Overall, this pattern indicates that parents sensibly named objects when the child's spatial attention was directed to the target. Finally, by the method of first and last reliable pairwise differences, the overlap with the

image center was reliably different for the target versus competitors at 4 s prior to the naming event and persisted until 1 s after the event for the naming events associated with learned object names and was reliably different for the target versus competitors at 3 s prior to the naming event and persisted until 1 s after the event for naming events associated with unlearned object names. The main results of the centering analyses are these: (1) The named target showed a clear temporal profile in

Table 1 Results of the growth curve analysis for object image size (left section) and centering (right section)

	Object Image Size (% of Image)				Centering (Average Distance to Center) (% of Half-Diagonal)			
	Estimate	95% HPD Interval		$p <$	Estimate	95% HPD Interval		$p <$
		Lower	Upper			Lower	Upper	
Intercept	3.20	2.72	3.69	.001	52.9	46.7	58.8	.001
Time	Not included in best-fit model				-1.7	-3.2	-0.1	.05
Time ²	-0.58	-0.74	-0.41	.001	4.1	2.5	5.6	.001
Object (target)	0.51	0.43	0.59	.001	-1.6	-2.3	-0.8	.001
Learning (learned)	0.08	-0.01	0.17	.093	-1.8	-2.7	-0.9	.001
Object (target) × learning (learned)	0.92	0.75	1.08	.001	-2.7	-4.2	-1.2	.001

Note. B -parameter estimates are presented with its respective 95 % highest posterior density interval and p -value calculated using a Markov chain Monte Carlo method; when meaningful, the category the estimate refers to is in parentheses

which the named target—but not the competitors—was increasingly more centered in the child’s view prior to the naming event and that this centering declined after naming, and (2) naming events associated with learned showed a higher centering advantage of the named target over the competitors than did the unlearned named targets.

The joint consideration of both the image size and centering analyses yields the following conclusion: Both centering and image size were dynamically related to the naming of an object by a parent and indicated that parents named objects when the target was being attended to by the child. However, learning also depended on the sustained visual dominance, as measured by image size and centering, of the named target over competitors.

It is likely that the two visual measures are not orthogonal but are codependent in a context of free-flow interaction. For example, a child or parent holding of an object so that the child is actively examining it during naming could bring the object closer to the child’s view, with the result of both a larger and more centered image of the object in the head camera. To determine the degree to which these two measures might be dynamically linked and, thus, redundant measures of the very same visual event, we repeated the GCA analysis by partialling out the effect of the second measure. Specifically we estimated the parameters of two models: the best-fit model structure for image size and for centering, but with the residuals of image size predicted by centering (using a linear regression), and the residuals of centering predicted by image size as the dependent variable.

In summary, this analysis revealed that although moderately correlated, $r = .48$, $p < .001$, image size and centering are not entirely redundant. The parameters that remained significant were the named object fixed effect, and the learning \times named object two-way interaction, when predicting residuals of image size ($p < .001$), and the quadratic time term, learning and named object fixed effects, and the learning \times named object two-way interaction when predicting residuals of centering ($p < .05$). Comparing these findings with the main results in Table 1, this analysis yielded the same general conclusion: Object size and centering of the target relative to competitors distinguished naming events associated with learned object names from those associated with unlearned object names, the visual dominance effect of named target versus competitors, and the higher visual. The sole qualitative difference was in the time terms, perhaps reflecting the similarities in the temporal pattern of both measures (a U-shaped curve that peaks at the naming event). This overall pattern suggests that object size and centering, although likely interdependent visually and in the sensory–motor aspects of the interaction that give rise to them, are also somewhat separable in their effects on learning and also, perhaps, in the specific behaviors by parents and infants that give rise to them.

Finally, to ensure that these conclusions did not depend on averaging the image sizes and centering of the competitors, a third set of analyses used the maximal value of the two competitors, rather than the mean; these analyses revealed the same basic findings.

Discussion

The results reveal the properties of visually optimal moments for toddlers to learn an object name: when the named object is visually larger and more centered than competitors and when that visual advantage is sustained for several seconds before and also after the naming event. The results are correlational and, as such, cannot specify the factors that created the observed visual signature for learned object names or the mechanisms through which limited visual competition and sustained attention benefit learning. However, the findings suggest that the sensory properties of naming moments matter. They also provide new insights into the assumptions about ambiguity in the input and also raise new hypotheses, at the visual level, about the specific challenges posed by scenes with multiple objects.

Contemporary theories of early object name learning begin with the problem of referential ambiguity and offer cognitive solutions to that problem: the inference of a speaker’s intended referent from social cues (e.g., Baldwin, 1995), the use of linguistic cues and innate biases (e.g., Lidz et al., 2003), and powerful statistical learning mechanisms (e.g., Xu & Tenenbaum, 2007). However, the present results tell us that for young learners, there is sometimes *little* ambiguity and that these moments of minimal visual ambiguity are strongly associated with object name learning. Not all naming moments had this property; many naming events associated with unlearned names were associated with multiple and nearly equal competitors for that name. Thus, the present results affirm the ambiguity often assumed and show that it also characterizes the visual level and the first-person view; and the results show that such ambiguity does make learning more difficult. But they also show there are very clean sensory moments when no additional cognitive processes would seem to be needed to determine the relevant object; no cognitive processes are needed because there is a sustained view in which just one object is much more salient in image size and centering than are possible competitors. One might conclude from these findings that there is no need to propose higher cognitive learning mechanisms, since young word learners might learn words only when there is minimal ambiguity at the visual level. Alternatively, these visually optimal moments may play a bootstrapping role, helping the child acquire or tune more cognitive and inferential processes that can succeed even given noisy input.

The dynamic visual properties of naming events associated with learning versus not learning the object name also suggest that there are *visual* limits on object name learning. This is a perspective that has not been considered in previous research but that is critical to understanding the mechanisms that underlie early object name learning and the properties of the learning environment that matter. Previous studies of adult visual processing show that multiple objects that are visually close to each other perturb both visual selection and representation in adults (e.g., Henderson, Chanceaux, & Smith, 2009). Recent studies suggest that the negative effects of clutter and crowding may be even more pervasive in toddlers (Oakes, Hurley, Ross-Sheehy, & Luck, 2010). Movement and change in the visual field can mandatorily capture attention in adults (see Knudsen, 2007) and also in toddlers (Columbo, 2001). Clearly, we need to understand these visual limits on early object name learning in greater detail. Indeed, the key factors in parent–child interactions with respect to early object name learning may be in limiting visual clutter and in sustaining selective attention on one object. Infant behavior itself may matter, since previous head camera studies suggest that views in which one object dominates are often linked to the toddlers' holding of the object (Smith et al., 2011; Yu et al., 2009). A large literature also suggests an important role for parent behavior, both as a top-down cue to attention (e.g., Tomasello & Akhtar, 1995) and also in terms of behaviors—holding, moving, and gesturing—that may directly structure the visual input. The present findings also suggest clear limits on what parents can do: Parents named objects when their infants' heads were spatially directed to the object (and the object was close to the child and centered in the view), and sometimes infants learned and sometimes they did not. Parents sometimes named the object when one object was visually dominant over competitors (and their infants learned), but they also sometimes named the object when the target and competitors were more equal in visual size (and their infants did not learn). This suggests that the child's view and its properties at the sensory level are not completely transparent to parents. Detailing the role of parent behavior and child behavior in structuring the bottom-up information and parent sensitivity to that information is a key issue for future research.

One potentially important finding with respect to the mechanisms underlying early word learning is the temporal duration of the visual advantage of the target over competitors for naming events associated with learning: beginning 6 s prior to the naming event and lasting 5 s after. This long duration could be indicative of the kind of factors—child activity and interest, parent activity in structuring the learning moment—that create optimal visual moments for learning object names and need not be essential to the mechanisms of learning. However, the increased stickiness of attention over time has been hypothesized to be important for sustained attention in toddlers (e.g., Richards & Cronise, 2000). Alternatively, the internal processes

that bind a name to an object may themselves take time and might, for example, require the formation of a stable visual representation of the object (Fennell, 2011; Ramscar, Yarlett, Dye, Denny, & Thorpe, 2010) prior to the naming event and/or maintenance of that visual representation (without replacement by another attended object) for some time after the heard name. These are hypotheses that need to be experimentally evaluated. In summary, the duration of sustained visual dominance of the target over the competitor observed in the present results may provide important clues as to how these optimal visual moments were created and also the mechanisms through which they benefit object name learning.

In conclusion, some early naming events are not ambiguous, not from the learner's view since there is but one dominant object in view. These may be optimal visual moments for mapping a name to an object and play a particular critical role for very young word learners. The differences in the visual properties of naming events associated with learned and unlearned object names also suggest potential visual limits on learning—in terms of clutter and in terms of sustained selective attention that endures over several seconds, limits that merit detailed experimental study.

Acknowledgments We thank Charlotte Wozniak, Amanda Favata, Amara Stuehling, and Andrew Filipowicz for collection of the data and Thomas Smith for developing data management and preprocessing software. This research was supported by National Science Foundation Grant 0924248, AFOSR FA9550-09-1-0665, NICHD grant R01HD 28675, and Portuguese Ministry of Education and Science Postdoctoral fellowship SFRH/BPD/70122/2010 awarded to Alfredo F. Pereira.

References

- Adolph, K. E., & Berger, S. E. (2006). Motor development. In W. Damon & R. Lerner (Eds.), *Handbook of child psychology: Vol 2* (6th ed., pp. 161–213). New York: Wiley.
- Allopenna, P. D., Magnuson, J. S., & Tanenhaus, M. K. (1998). Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory and Language*, 38(4), 419–439.
- Aslin, R. N. (2009). How infants view natural scenes gathered from a head-mounted camera. *Optometry and Vision Science: American Academy of Optometry*, 86(6), 561.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of memory and language*, 59(4), 390–412.
- Baldwin, D. A. (1995). Understanding the link between joint attention and language. In: *Joint attention: Its origins and role in development*, ed. C. Moore & P. J. Dunham. Erlbaum, 131–158.
- Bates, D. (2005). Fitting linear mixed models in R. *R news*, 5(1), 27–30.
- Bates, D. (2012). Linear mixed model implementation in lme4. *Ms., University of Wisconsin*.
- Bertenthal, B., & Von Hofsten, C. (1998). Eye, head and trunk control: The foundation for manual development. *Neuroscience & Biobehavioral Reviews*, 22(4), 515–520.
- Booth, A. E., & Waxman, S. R. (2009). A horse of a different color: Specifying with precision infants' mapping of novel nouns and adjectives. *Child Development*, 80, 15–22.

- Cicchino, J. B., Aslin, R. N., & Rakison, D. H. (2010). Correspondences between what infants see and know about causal and self-propelled motion. *Cognition*, *118*, 171–192.
- Colombo, J. (2001). The development of visual attention in infancy. *Annual Review of Psychology*, *52*(1), 337–367.
- Fennell, C. T. (2011). Object familiarity enhances infants' use of phonetic detail in novel words. *Infancy*
- Frank, M. C., Goodman, N. D., & Tenenbaum, J. B. (2009). Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science*, *20*(5), 578–585.
- Henderson, J. M., Chanceaux, M., & Smith, T. J. (2009). The influence of clutter on real-world scene search: Evidence from search efficiency and eye movements. *Journal of Vision*, *9*(1).
- Iverson, J. M. (2010). Developing language in a developing body: The relationship between motor development and language development. *Journal of Child Language*, *37*(02), 229–261.
- Kannass, K. N., Oakes, L. M., & Shaddy, D. J. (2006). A longitudinal investigation of the development of attention and distractibility. *Journal of Cognition and Development*, *7*(3), 381–409.
- Knudsen, E. I. (2007). Fundamental components of attention. *Annu. Rev. Neurosci.*, *30*, 57–78.
- Lawson, K. R., & Ruff, H. A. (2004). Early focused attention predicts outcome for children born prematurely. *Journal of Developmental & Behavioral Pediatrics*, *25*(6), 399–406.
- Lidz, J., Waxman, S. R., & Freedman, J. (2003). What infants know about syntax but couldn't have learned: Experimental evidence for syntactic structure at 18 months. *Cognition*, *89*, 295–303.
- Mayer, D. L., & Fulton, A. N. (1993). Development of the human visual field. In K. Simons (Ed.), *Early Visual Development, Normal and Abnormal* (pp. 117–129). New York, NY: Oxford University Press.
- Mirman, D., Dixon, J. A., & Magnuson, J. S. (2008). Statistical and computational models of the visual world paradigm: Growth curves and individual differences. *Journal of memory and language*, *59*(4), 475–494.
- Oakes, L. M., Hurley, K. B., Ross-Sheehy, S., & Luck, S. J. (2010). Developmental changes in infants' visual short-term memory for location. *Cognition*, *118*, 293–305.
- R development core team. (2008). R: A language and environment for statistical computing. *R Foundation Statistical Computing*.
- Ramscar, M., Yarlett, D., Dye, M., Denny, K., & Thorpe, K. (2010). The effects of feature label order and their implications for symbolic learning. *Cognitive Science*, *34*(6), 909–957.
- Richards, J. E., & Cronise, K. (2000). Extended visual fixation in the early preschool years: Look duration, heart rate changes, and attentional inertia. *Child Development*, *71*(3), 602–620.
- Smith, L. B., Yu, C., & Pereira, A. F. (2011). Not your mother's view: The dynamics of toddler visual experience. *Developmental Science*, *14*(1), 9–17.
- Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, *106*, 1558–1568.
- Tomasello, M., & Akhtar, N. (1995). Two-year-olds use pragmatic cues to differentiate reference to objects and actions. *Cognitive Development*, *10*(2), 201–224.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*, 245–272.
- Yoshida, H., & Smith, L. B. (2008). What's in view for toddlers? using a head camera to study visual experience. *Infancy*, *13*(3), 229–248.
- Yu, C., Smith, L. B., Shen, H., Pereira, A. F., & Smith, T. (2009). Active information selection: Visual attention through the hands. *Autonomous Mental Development, IEEE Transactions on*, *1*(2), 141–151.