Moving beyond “nouns in the lab”: Using naturalistic data to understand why infants’ first words include *uh-oh* and *hi*

Kennedy Casey^{1,2}, Christine E. Potter^{1,3}, Casey Lew-Williams¹, & Erica H. Wojcik⁴

¹ Princeton University

² University of Chicago

³ University of Texas at El Paso

⁴ Skidmore College

Correspondence

Kennedy Casey

Department of Psychology, Princeton University

Princeton, NJ 08540, USA

Email: kcasey@alumni.princeton.edu

Data Availability

The data that support the findings of this study are available at <https://osf.io/tdbqn/>

Funding Information

National Institute of Child Health and Human Development, Grant/Award Number: F32HD093139 and R01HD095912; Princeton University

Conflict of Interest

The authors declare that there is no conflict of interest.

Ethics Approval

All experimental protocols, including procedures for obtaining informed consent, were approved by the Princeton University Institutional Review Board.

Acknowledgements

This work was supported by Princeton University (fellowship awarded to KC) and by the National Institute of Child Health and Human Development (F32HD093139 to CEP, R01HD095912 to CLW). We thank the participating families and the members of the Princeton Baby Lab, especially Matthew Weatherhead and Ellie Breitfeld.

Research Highlights

- One-year-old infants were tested on their comprehension of early-learned, routine-based ‘everyday words’ (e.g., *uh-oh*, *hi*, *more*, *up*, and *all-gone*).
- Standard, lab-based eye-tracking measures failed to reveal evidence of reliable comprehension of everyday words.
- Corpus-based observational analyses found that everyday words co-occur with highly variable visual referents in real-world learning environments.
- Everyday words appear in more consistent situational contexts, and stability at this level of the input may be linked to early learnability.

Abstract

Why do infants learn some words earlier than others? To explain how and when words are learned, existing theories of word learning prioritize visual information and draw mainly on lab-based studies of noun-to-object mapping. However, words that are more abstract than object nouns, such as *uh-oh*, *hi*, *more*, *up*, and *all-gone*, are typically among the first to appear in infants’ vocabularies. We combined a behavioral experiment with naturalistic observational research to explore how infants learn and represent this understudied category of high-frequency, routine-based non-nouns, which we term ‘everyday words.’ In Study 1, we found that conventional eye-tracking measures of comprehension were insufficient to capture 10- to 16-month-old infants’ emerging understanding of everyday words. In Study 2, we analyzed the visual and social scenes surrounding caregivers’ and infants’ use of everyday words in a naturalistic video corpus. This ecologically-motivated research revealed that everyday words rarely co-occurred with consistent visual referents, making their early learnability difficult to reconcile with dominant word learning theories. Our findings instead point to complex patterns in the types of situations associated with everyday words that could contribute to their early representation in infants’ vocabularies. By leveraging both experimental and observational methods, this investigation underscores the value of using naturalistic data to broaden theories of early learning.

Keywords: word learning; language input; naturalistic recordings; ecological validity; eye tracking

Moving beyond “nouns in the lab”: Using naturalistic data to understand why infants’ first words include *uh-oh* and *hi*

Early word knowledge is often viewed as a difficult mapping problem. For instance, a child must determine that *cup* refers to the object that holds their juice and does not refer to their spoon, their kitchen table, or other neighboring objects in their home. Impressively, by nine months, infants reliably show evidence that they understand several common concrete nouns, including names for objects that they see and interact with frequently (Bergelson & Swingley, 2012, 2015; Kartushina & Mayor, 2019; Parise & Csibra, 2012; Tincoff & Jusczyk, 1999, 2012). Various proposals have been put forth to explain such learning, including those that emphasize infants’ use of social-attentional cues, such as speakers’ eye gaze, to determine intended object referents (Çetinçelik et al., 2021) and infants’ tracking of word-object co-occurrences (Smith & Yu, 2008; Stevens et al., 2017). However, existing accounts of word learning primarily focus on infants’ ability to map nouns onto concrete objects and thus address only a subset of the words that young children must learn.

‘Everyday words’: An understudied subset of infants’ early vocabularies

Not all of infants’ early words are object names. Young children learn verbs, adjectives, and sounds, along with other words that do not tidily fit into established lexical categories. In fact, *uh-oh* is the fourth most commonly produced English word at 16 months, preceded only by *mommy*, *daddy*, and *ball* (Frank et al., 2017). Other non-nouns, such as *hi*, *more*, *up*, and *all-gone*, occur frequently in infant-directed speech and appear consistently in infants’ earliest productions across languages, as observed in prior research using diverse methods, such as diary studies, parent surveys, and naturalistic observations (e.g., Bates et al., 1994; Bloom et al., 1993;

Bowerman, 1978; Caselli et al., 1995; Frank et al., 2021; Gleason et al., 1984; Lieven et al., 1992; Nelson, 1981; Tardif et al., 2008; Tomasello, 1987). These routine-based, social words (hereafter—‘everyday words’) make up approximately 25% of the first 50 words in English and 25% of the first 10 words across 15 languages (Frank et al., 2017, 2021). Despite this prevalence, these routine-based words remain not only underexplored in developmental research but virtually unacknowledged in theories of early word learning. This oversight raises the question of whether current proposals of learning mechanisms, generated primarily from lab-based studies on the learning of concrete nouns, can adequately explain early word learning.

‘Nouns in the lab’: The standard approach to assessing early comprehension via eye tracking

The tendency of word learning theories to overlook everyday words could reflect entrenched assumptions about nouns holding a privileged position in the early lexicon, but it also may simply reflect the difficulty in visually depicting their meanings—a necessity for standard lab-based comprehension tasks. That is, what exactly does *uh-oh* look like? Language scientists have largely depended on eye-tracking studies to assess young children’s recognition of familiar words. For example, in the looking-while-listening procedure (LWL, Fernald et al., 2008), infants are typically presented with two side-by-side images of common objects or animals, and researchers measure their ability to look at the appropriate image when it is labeled. Looking time to the correct image serves as a proxy for comprehension because children tend to look more reliably at the target referent as their lexical knowledge develops (Fernald et al., 2006). This method critically depends on researchers’ ability to depict imageable, common referents for the target words, making this approach particularly well-suited for testing comprehension of concrete nouns with object-based meanings.

The LWL procedure has also been applied to assess young children's understanding of imageable words other than nouns, such as familiar verbs (Golinkoff et al., 1987; Valteau et al., 2018) and adjectives (Fernald et al., 2010; Forbes & Plunkett, 2019). Two studies have included a few everyday words when testing comprehension of non-nouns (e.g., verbs, adjectives, proper names) and yielded some weak evidence of comprehension in 9- to 16-month-olds (Bergelson & Swingley, 2013; Syrnyk & Meints, 2017). However, infants' comprehension of everyday words has yet to be systematically examined.

The current studies: Everyday words in the lab and in the world

To address this gap in our knowledge about infants' understanding of early-learned everyday words, we first conducted an in-lab study. Following current conventions and best practices, we tested infants' comprehension of everyday words with an eye-tracking paradigm and pre-registered our design and planned analyses (<https://osf.io/z5qxf>). However, unlike studies of early-learned nouns, our study failed to reveal reliable evidence of everyday word understanding. We report the details of this unsuccessful behavioral experiment and consider why standard eye-tracking measures did not allow infants to demonstrate their emerging knowledge.

Motivated by the lack of evidence for comprehension in Study 1, we took a different approach to exploring infants' representations of everyday words in Study 2. Specifically, we conducted detailed descriptive analyses of a naturalistic video corpus. In doing so, we captured novel characteristics of the real-world visual and social contexts that surround everyday words in the lives of infants. The results of this ecologically-motivated work challenge long-held assumptions about the processes by which early word meanings are learned and introduces theoretical and methodological considerations for future research.

Study 1: Behavioral Experiment

Using a standard LWL paradigm, we tested whether infants could associate each of 12 everyday words with a corresponding referent (e.g., *uh-oh*: fallen cup; *hi*: person waving; Figure 1). Since the target words are highly frequent in infant-directed language (MacWhinney, 2000) and are reported to be understood, on average, by 60.5% of English-learning 12-month-olds (Frank et al., 2017) as well as 69.7% of the infants in our sample (see Method), we predicted that conventional looking-time measures would provide the first experimental evidence of early comprehension of everyday words.

Method

Participants

Participants were 33 full-term, monolingual English-learning infants (13 female), ranging in age from 10 to 16 months ($M = 13.2$ months). All infants had no reported hearing or vision impairments and were exposed to English at least 85% of the time. Parents provided informed consent, and families received a small gift in exchange for their participation. We pre-registered a sample size of 50 to match that of a similar previous study (Bergelson & Swingley, 2013). However, data collection was prematurely stopped due to COVID-19. At this point, 52 infants had participated but 19 did not meet inclusion criteria for analysis based on pre-registered guidelines due to fussiness ($n = 9$), equipment malfunction ($n = 1$), or failure to contribute data on at least 50% of test trials ($n = 9$).

Stimuli

All experimental stimuli are available on OSF (<https://osf.io/tdbqn/>). Infants were presented with 24 test images (two per target word), organized into yoked pairs (*uh-oh/hi*, *wow/all-gone*, *bye-bye/yum*, *no/night-night*, *up/more*, *shh/thank-you*). Target words were determined based on CDI production norms at 16 months (Frank et al., 2017) as well as frequency counts from the Child Language Data Exchange system corpus (CHILDES, MacWhinney, 2000) for children up to 17 months of age. The final 12 items were selected among candidate words from these two sources based on researchers' intuitions about imageability (see Table S1 for stimuli characteristics).

Following established traditions for designing LWL experiments, visual referents were chosen based on researchers' views about what infants around 12 months of age typically see when they hear the target words produced in naturalistic contexts. Target images were informally matched for visual salience and yoked such that their referent scenes would not be easily confusable (e.g., *bye-bye* was not matched to *hi* since both involve a person waving).

Auditory stimuli consisted of natural recordings of the 12 target words (two tokens per word), produced in infant-directed speech by a female native English speaker and left unedited to preserve typical intonation. Target words ranged from 410ms to 1240ms in duration ($M = 820\text{ms}$) and were normed to a standard mean intensity of 65dB (Boersma & Weenink, 2016).

Figure 1. Sample target images (full set available on OSF). Target images were reported by caregivers to generally match infants' typical visual experience. The mean image familiarity rating was 2.9/4 across all participants and 3.1/4 for the subset of participants with reported target word comprehension.

Procedure

Participants sat on their parents' lap in a dimly-lit testing room and viewed two images at a time on a 55-inch TV monitor display. Caregivers wore opaque sunglasses and were asked to avoid directing their child's attention during the study. All 12 target words were tested twice, each time with a different target image, resulting in 24 test trials. Trial order was pseudorandomized such that there were at least two trials between repetitions of a word ($M = 11$ trials between repetitions), and target side was counterbalanced to ensure an equal number of left and right target trials across the study. Participants were randomly assigned to one of two trial orders, where the second was the reverse of the first. Infants first saw two practice trials with

common concrete nouns as targets (*ball* and *dog*). On each of the following test trials, images appeared on the screen for 2500ms before the onset of the auditory stimulus. Then, participants heard two repetitions of the target word in isolation, with a fixed 2200ms delay between the onset of the first and second repetition of the word. After the second repetition, images remained on the screen for an additional 2000ms, resulting in test trials that were 6900ms long (+ 500ms inter-trial interval). Infant-friendly, nonverbal videos or verbal reinforcement appeared between blocks of four test trials to maintain engagement.

After the study, parents completed a vocabulary and image familiarity survey, where they were asked to report whether their child understood and/or said each of the 12 target words and to rate the visual stimuli for familiarity. Results from the vocabulary checklist indicated that all 12 everyday words were understood by a majority of participants. On average, 69.7% of infants reportedly understood ($SD = 13.3\%$, range = 45.5-84.8%) and 14.4% reportedly produced each target word ($SD = 10.1\%$, range = 3.0-30.3%). Parents attributed a relatively high degree of familiarity to our target images, on average, meaning that the chosen scenes were thought to match infants' typical visual experience. The mean image familiarity rating across words was 2.9/4 overall (median = 3.0, $SD = 1.0$, range = 0.5-4) and 3.1/4 for the subset of infants with reported target word comprehension (median = 3.0, $SD = 0.89$, range = 1.0-4). Additionally, children's global receptive and expressive vocabulary knowledge was assessed using the CDI: Words & Gestures (Fenson et al., 1994), but we found no relationship between vocabulary size and LWL task performance (see SI for details).

Infants contributed usable data on 19.4 out of 24 test trials, on average ($SD = 3.4$ trials). Videos of participants' eye movements were coded offline, frame by frame, in 33-ms intervals. For each frame, a trained coder, blind to target condition, determined whether the infant was

looking at the left image, right image, away from both images, or shifting between images.

Following our pre-registered plan, trials were excluded if participants looked away from both images for 50% or more of the critical window (367 to 4000ms following target word onset; consistent with Bergelson & Swingley, 2013). To ensure reliability in gaze coding, 20% of videos were re-coded by a second coder. Inter-coder reliability was high, with coders agreeing on infants' gaze location within a single frame on 99% of frames overall, and for a more conservative measure, on 97% of frames surrounding shift events.

Results

Overall performance on the LWL task was at chance level ($M = 0.51$, $p = 0.96$, Wilcoxon test). That is, infants did not look reliably more to labeled target referents, relative to unlabeled distracters, during the critical window from 367 to 4000ms following target word onset. To better understand the dynamics of looking behavior, as planned in our pre-registration, we analyzed the time course data using growth curve analysis (GCA, Mirman, 2014, see SI for model details). Infants' accuracy in looking to target images was above chance based on the intercept term ($Estimate = 0.52$, $SE = 0.06$, $t = 8.29$, $p < 0.001$; Figure 2A). However, both the linear and quadratic time terms were not significant, suggesting that participants' looks to the target image did not increase significantly after target word onset ($Estimate = 0.01$, $SE = 0.19$, $t = 0.03$, $p = 0.98$) and looking behavior did not follow a parabolic trajectory ($Estimate = 0.11$, $SE = 0.12$, $t = 0.89$, $p = 0.38$). While the significant intercept term points to some weak evidence of everyday word comprehension, the profile of target looking in the current study does not match the profile that is seen for early-learned concrete nouns at the same age (e.g., Bergelson & Swingley, 2015). Moreover, exploratory analyses of item effects (Figure 2B; Figures S1-S2) as well as age and

image familiarity effects found no reliable evidence of comprehension of everyday words (Tables S2-S3).

Figure 2. (A) Growth curve estimates of empirical logit-transformed mean proportion of target looking during the critical window, offset by 0.5. Points reflect means over participants and items for each 33-ms time bin in the critical window. (B) Salience-corrected mean accuracy across item-pairs. Accuracy was calculated as the proportion of time spent looking at a given test image when it appeared as the labeled target image, relative to the proportion of time spent looking at the same image when it appeared as the unlabeled distracter. Points reflect means over participants. Chance-level performance (0.0) is indicated by the dashed line. Error bars reflect standard errors of the mean, and asterisks denote above-chance performance after correcting for multiple comparisons ($p < 0.008$).

Discussion

Results of Study 1 revealed that, on average, infants did not significantly increase their looking to the target image after word onset. Conventional looking-time measures provided little to no evidence of comprehension even though the tested words were reported to be understood by

a majority of participants and the test images were considered by caregivers to match infants' typical visual experience. These null effects could be due to a true lack of understanding of everyday words, but that possibility conflicts with evidence of early comprehension and production obtained from thousands of parent-report surveys (Frank et al., 2017).

Because caregivers reliably report that infants, both generally and in the current study, do understand everyday words, infants' failure in the behavioral experiment could be explained by stimulus limitations stemming from poor adult intuitions about infants' representations of these words. Although our visual stimuli for Study 1 were validated by parent report, it is nevertheless possible that adults' assumptions about common referents failed to resemble infants' typical visual experience associated with these words (i.e., *uh-oh* might not regularly occur in the presence of a fallen cup). This mismatch is particularly possible given the routine-based nature of everyday words. Whereas the experimenter-intuited referents of early-learned concrete nouns (e.g., *cup* or *ball*) are likely to match the common tangible objects seen by infants when these words are uttered (Custode & Tamis-LeMonda, 2020), the meanings of everyday words may not be as adequately expressed by unimodal, static images chosen by developmental scientists.

Study 2

Current theories of word learning primarily explain infants' learning of concrete nouns with stable visual referents and are supported by evidence of robust noun-object co-occurrence in naturalistic word learning environments (Custode & Tamis-LeMonda, 2020; Pereira et al., 2014). However, without knowing the characteristics of the input surrounding everyday words, we cannot yet determine whether the same principles used to account for the learning of concrete

nouns also extend to other early-learned words, including everyday words¹. One prior study found that some abstract words, including a few routine-based words like *uh-oh* and *bye-bye*, were less likely than concrete nouns to be uttered in the presence of their referents (Bergelson & Swingley, 2013), but these data only describe whether or not the experimenter-intuited referent was present when the word was uttered. This top-down approach, also used in Study 1, contrasts with the primarily bottom-up approach used in the next part of our investigation.

Study 2 bypassed adult assumptions about word meanings by providing a detailed descriptive analysis of the input surrounding everyday words in longitudinal, at-home recordings of five one-year-olds from the Providence corpus (Demuth et al., 2006). A few prior studies have described examples of some early meanings of everyday words (e.g., *up* referring to a child being picked up, Tomasello, 1987; or *bye-bye* being associated with people leaving, Gleason & Weintraub, 1976). The primary aim of our naturalistic observational research was to describe all attested meanings in children’s input and thereby obtain a comprehensive picture of the possible referents associated with everyday words.

We first quantified the consistency with which everyday words co-occurred with different visual referents. However, after quickly discovering the vast variability of the visual referents associated with everyday words in a first coding pass, we also defined ‘situational contexts’—broader categories of routine-based usage that collapse across basic visual features of referent scenes (e.g., the exact objects present) to acknowledge and describe some level of consistency in

¹ There is an existing disagreement surrounding whether more abstract, routine-based words (e.g., *hi*, *bye-bye*) are referential words (e.g., Gleason & Weintraub, 1976; Lieven et al., 1992). Even if everyday words are interpreted as a qualitatively different verbal signal, description of the input surrounding production is still critical for determining *how* they become different words from the perspective of the learner.

children's input. As an example, if the word *uh-oh* was produced when a child dropped their cup, the visual description would be 'cup falling,' and the situational description would be 'object falling.' A secondary aim was to assess the validity of our approach in Study 1, and we did this with respect to both the visual and situational appropriateness of our stimuli. By describing the characteristics of infants' real-world input, we can evaluate whether current theories of word learning can explain infants' early learning of everyday words.

Method

Coding Procedure

Recording sessions (114 total) were coded for five children (range = 11 to 24 months) in the Providence corpus (Demuth et al., 2006). All utterances containing one of the 12 target everyday words were found using the PhonBank online database (Rose & MacWhinney, 2014). We watched the video accompanying each utterance to determine the visual scene and usage context of the target word. As needed, we relied on information from preceding or following utterances to better understand the intended usage. Tokens were marked as 'indeterminable' and excluded from analyses if the usage could not be reasonably inferred from the available media (5.85% of the dataset), most often due to the speaker/target child being out of frame during the moment of target word production.

We extracted 11,920 total tokens across all 12 target everyday words (Table S4). Frequency varied considerably across items and children, likely reflecting natural variation along with sampling biases. For instance, *night-night* had the lowest token count, at least in part because recording sessions primarily took place during daytime hours, and counts for *yum* varied drastically across children because some families did not record during mealtime. Crucially,

though, even for low-frequency items, we extracted over 150 total tokens, and 11 of 12 items included tokens sampled from all five households ($M = 200.2$, range = 30.6-633.8 tokens per household).

For each token, we annotated eight variables of interest, including (1) the speaker who produced the target word, (2) whether the speech was directed to the target child or overheard, (3) whether the speaker (if not the target child) was present, (4) whether the target child was present, (5) the visual referent, (6) whether the token was a visual match to our experimental stimuli, (7) the situational context, and (8) whether the token was a situational match to our experimental stimuli. Visual referents were coded as the exact objects, actions, and/or people involved in the referent scene (e.g., *uh-oh*: cup falling, block tower falling, child falling, crayon breaking, phone ringing, etc.). Because poor video quality sometimes compromised our ability to determine the speaker/target child's view when the word was uttered, our coding reflects the 'idealized' referent for each token, determined based on a combination of the available video and audio. For instance, the referent of *uh-oh* could be coded as 'cup falling' even if the video did not capture the dropping event or if there was temporal displacement between the dropping event and utterance of the target word, so long as the audio unambiguously reflected this intended usage. Situational contexts refer to broader clusters of usage that describe, at a level higher than the specific visual features of the scene, the intended meaning of the target word (e.g., *uh-oh*: any object falling). Situational context categories were generated by grouping together similar referents from the visual-level coding (e.g., *uh-oh*: cup falling and block tower falling were coded as the same situational context since both involved an object falling; see Table 2 for the top three situational contexts by item). To evaluate the validity of our experimental stimuli, we also characterized referent scenes according to whether they matched the visual referents and/or situational contexts

depicted in the test images from Study 1 (e.g., *uh-oh*: cup falling = ‘visual match,’ any object falling = ‘situational match’).

To ensure reliability in descriptive coding, a randomly selected 5 tokens per target word per child were re-coded by a second coder. High reliability was achieved through an iterative training process. This process involved discussion of discrepancies between the primary and naïve coders on a separate training set (5 tokens per target word per child), with initial estimates of inter-coder reliability at 84%. Following training, coders agreed on 91% of decisions.

Results

Everyday words do not reliably co-occur with consistent visual referents We first determined the frequency with which words appeared with different visual referents (e.g., *uh-oh*: cup falling, block tower falling, child falling, crayon breaking, etc.). Our analyses revealed that the visual input surrounding everyday words exhibits high variability within and across households, over time, and in both infants’ and adults’ productions.

We found that the most common visual referent for each everyday word in the naturalistic corpus (e.g., child falling for *uh-oh*, cereal for *more*, or play doh for *wow*; Table 1) captured a small fraction of the input. The top visual referent co-occurred with only 9.0% of word tokens on average (range = 1.4-20.4%; Figure 3A) and never matched for all five children. In fact, out of 4,120 total unique visual referents (summed across all words), only 13 appeared at least once for all children in the corpus (Table S5). Although this cross-household variability stems in part from sampling limitations of the corpus, it suggests that individual children’s early representations are idiosyncratic and underscores the visual inconsistency of everyday words, which is reflected in children’s productions (Figure S9) and persists across developmental time (Figures S13-15).

Table 1. Top referents by target word

Target word	Most common referent	2nd most common referent	3rd most common referent
<i>all-gone</i>	crayons gone	done reading	cookies gone
<i>bye-bye</i>	dad leaving	camera (done recording)	ending phone call
<i>hi</i>	mom greeting child	answering phone	answering pretend phone call
<i>more</i>	more cereal	more beans	more cheese
<i>night-night</i>	goodnight book	animals going to bed in book	eyes closing on flap book
<i>no</i>	child not wanting to read	child not wanting to share toys with mom	child not wanting specific book
<i>shh</i>	wheels on the bus song - parents on the bus say <i>shh</i>	hush song - little babies love to sleep	child screaming
<i>thank-you</i>	child giving mom play doh	child giving mom block	child giving mom bubble wand
<i>uh-oh</i>		microphone falling off	broken crayon
<i>up</i>	child falling	pop up toy	pop up book
<i>wow</i>	mom picking child up	bubbles	child falling
<i>yum</i>	play doh	lunch (unspecified)	cheese
	yogurt

Figure 3. (A) Distribution of top visual referents across target words. (B) Distribution of top situational contexts across target words. Dark gray bars refer to tokens not matching the top three visual referents, and light gray bars refer to tokens where intended usage was indeterminable (e.g., due to the speaker and target child being out of frame). The plot for visual referents shows particularly large variability beyond the top three referents, while the plot for situational contexts shows more stability at this level of the input.

Everyday words appear in diverse situational contexts, but there are one or more dominant contexts for each word

In a second set of analyses, we determined the frequency with which everyday words appeared with different situational contexts (e.g., *uh-oh*: any object falling, any person falling, any object breaking). It is possible that despite being visually inconsistent, the input surrounding everyday words is more stable when considering the general context of the event-based scene. Although we still found vast variability at this broader level of the input, everyday words co-occurred more reliably with situational contexts than with visual referents. On average, the most common single situation accounted for 48.1% of the input (range = 12.8-94.0%; Figure 3B). The top three situational contexts captured a majority of tokens for 10 of 12 items, with proportions ranging from 31.5 to 100.0% of tokens ($M = 72.4\%$). This pattern also appeared across children. The top situation accounted for 45.2% of each child's input, on average (range = 39.8-94.0%; Figure S12), and four of five children shared the same top situation, on average across words

(range = 2-5). The dominance of the top situations also held across time (Figures S13-S15) and was reflected in infants' productions (Figure S11). Thus, while the visual information associated with everyday words is highly variable, the situational-level input appears more stable, both within and between households and over time.

Table 2. Top situations by target word

Target word	Most common situation	2nd most common situation	3rd most common situation
<i>all-gone</i>	done playing	done eating	person leaving
<i>bye-bye</i>	person leaving	done interacting with animal	putting away other inanimate object
<i>hi</i>	initiating interaction with person	answering phone	initiating interaction with animal
<i>more</i>	eating/requesting more food	playing with/requesting more toys	requesting more of another inanimate object
<i>night-night</i>	going to bed	putting toy to sleep	pretending to sleep
<i>no</i>	not wanting to do something	doing something not allowed	not wanting someone else to do something
<i>shh</i>	someone sleeping	someone screaming	someone crying
<i>thank-you</i>	exchanging object	being helpful	following directions
<i>uh-oh</i>	object falling	person falling	object breaking
<i>up</i>	referring to location (e.g., "up here")	cleaning up	picking up person
<i>wow</i>	playing with toy	performing impressive action	referring to other inanimate object
<i>yum</i>	eating/referring to food	pretending to eat	N/A

Experimenter-intuited stimuli from Study 1 depicted plausible visual referents and generally matched dominant situational contexts

To assess the validity of our experimental stimuli, we determined the frequency with which the visual referents and situational contexts depicted in the images from Study 1 appeared in the naturalistic corpus. Unsurprisingly given the high degree of visual variability associated with everyday words, we found that visual matches to our experimental stimuli were rare across words ($M = 7.7\%$, range = 0.0-17.4%; Figure S5). However, the exact Study 1 visual referents appeared at least once for 11 of 12 everyday words, providing evidence that our experimenter-intuited scenes were plausible referents. Moreover, the top situations that emerged in the naturalistic corpus largely conformed with our broad intuitions about everyday word meanings. Situational matches to Study 1 stimuli were common ($M = 46.4\%$, range = 7.3-81.9%) and even mapped onto the top situational context for nine of 12 items. Together, these findings suggest that infants' failure to look to the appropriate images in Study 1 cannot necessarily be attributed to implausible test scenes and instead demonstrate that everyday words may not easily be captured by a single visual token, both within and across infants.

Exploratory analyses: Visual and situational stability may contribute to early learnability

Next, we explored how the consistency with which everyday words co-occurred with different visual referents and situational contexts may affect children's learning. That is, do children learn an everyday word earlier if it appears with consistent visual referents and situational contexts? We hypothesized that children whose visual and situational input for a given everyday word was more stable would be more likely to produce that word, relative to infants whose input was more variable.

We fit a linear mixed-effects model predicting children's rate of production of individual target everyday words, relative to their total rate of production, on the basis of visual and situational stability of the input (see SI for model details). Visual stability was calculated as the proportion of tokens corresponding to a repeated visual referent for each item and child (e.g., for *uh-oh*, what proportion of the time did each child hear this word in the presence of a visual referent that had appeared previously?). Situational stability was measured as the proportion of tokens mapping onto the top situational context for each item and child (e.g., for *uh-oh*, what proportion of the time did each child hear this word used to refer to any object falling?). We found that both visual ($Estimate = 1.11$, $SE = 0.44$, $t = 2.53$, $p = 0.01$) and situational ($Estimate = 0.72$, $SE = 0.35$, $t = 2.10$, $p = 0.04$) stability positively predicted production. Notably, visual and situational stability were not highly correlated. This result suggests that infants who heard a given word in the presence of more consistent visual referents and more consistent situational contexts were more likely to produce those words, relative to infants who experienced more variable visual and situational input. Thus, stability at multiple levels of children's input could support their learning.

Discussion

Results of Study 2 demonstrated that the visual input surrounding everyday words is highly variable. While dominant theories of word learning tend to rely on referent stability, we found that co-occurring visual referents varied both within and across children. This observation challenges long-held assumptions that imageable, concrete words are the first-learned (e.g., Gentner, 1982) because everyday words—despite their visual inconsistency—are typically among the first to appear in infants' vocabularies (Frank et al., 2017, 2021). These findings also call into question the validity of using eye tracking to assess infants' knowledge of less

imageable, routine-based words and therefore help explain the lack of reliable looking to target images in our behavioral experiment.

Our intuited visual referents from Study 1 appeared rarely in the naturalistic video corpus (i.e., ~8% of tokens). However, rather than reflecting invalid adult assumptions about early meanings, the low proportion of exact matches to our experimental stimuli underscores the visual inconsistency associated with everyday words. Even the most common visual referent captured only ~9% of the input and also varied across children, suggesting that the meanings of everyday words are unlikely to be fully represented by a single visual token or event. Notably, there is greater variability of visual scenes surrounding everyday words compared to the variability between exemplars within a concrete object category (Clerkin et al., 2017). While the distribution of different *cup* tokens, for example, is anchored on a frequent prototypical visual token and described by a function, the variability in everyday word visual referents has no obvious anchor or unifying structure. Thus, the characteristics of the visual input surrounding everyday words appear largely incompatible with current approaches to both testing and understanding early word learning.

Everyday words also appeared in diverse situational contexts, though the input was generally skewed toward one or more dominant contexts for each item. The top situation accounted for nearly half of the input, on average across words. That is, *uh-oh* does refer to things falling with some reliability. This finding suggests that even though the visual information associated with everyday words varies considerably, the situational information surrounding utterances appears more stable and—if children focus on this context level—could support word learning. Still, it is notable that the situational contexts associated with everyday words were far from uniform, particularly compared to the stability in the visual input surrounding concrete

nouns. While concrete nouns have been found to co-occur with the same referent over 80% of the time (Bergelson & Swingley, 2013; Custode & Tamis-LeMonda, 2020; Pereira et al., 2014), we found that everyday words co-occurred with the same situation only 48% of the time, indicating that infants must contend with substantial variability even at this broader level of the input to learn everyday words.

General Discussion

Using experimental and descriptive approaches, this investigation examined a common yet generally unexplored category of early-learned words, which we term everyday words. These high-frequency, routine-based words, such as *uh-oh* and *hi*, are well-represented in infants' earliest productive vocabularies, yet their learnability cannot be explained by current theories of word learning. Study 1 used eye-tracking measures to test whether one-year-old infants could associate everyday words with experimenter-intuited visual referents but failed to detect evidence of reliable comprehension. Study 2 revealed that the real-world visual input surrounding everyday words is highly unstable but that these words may appear in relatively more consistent situational contexts.

Limitations of lab-based measures of word comprehension

Early word knowledge encompasses more than just the ability to map labels onto stable, concrete referents. For example, infants must generalize words to new contexts and use their limited productive vocabularies to communicate their socio-emotional needs. Yet, lab-based studies generally define word comprehension in a rather restricted sense. By requiring that infants look at, or even point to, something to demonstrate understanding, standard behavioral measures may be insufficient to capture infants' emerging representations of words that may not have

direct visual associates (Wojcik et al., 2022). When presenting only two-dimensional images, or even videos (e.g., Bergelson & Swingley, 2013; Syrnyk & Meints, 2017) of intended referents, highly-controlled lab-based experiments generally lack the richness and diversity of infants' real-world language environments (Nastase et al., 2020; Reuter et al., 2021; Tamis-LeMonda et al., 2017) and may therefore eliminate connections to socio-emotional, pragmatic, or contextual information and other cues that are useful for word recognition. The low-dimensionality of common lab-based measures may be especially problematic for everyday words, which by their routine-based nature may be tied to less-visible situational factors, such as affective state (e.g., *uh-oh* is likely to be produced under more dramatic circumstances than many concrete nouns, Ponari et al., 2018) or event timing (e.g., *uh-oh* is likely to coincide with salient event transitions, such as when the child falls, *hi* is likely to occur when a new interaction begins, which creates opportunities for learning during moments of heightened attention, Kosie & Baldwin, 2019). Furthermore, early word knowledge may be highly idiosyncratic, as evidenced by the cross-household variability observed in our study, making it difficult to capture early comprehension with one standard set of stimuli. These methodological issues may be especially pronounced for everyday words.

The richness and complexity of learning contexts for everyday words

Naturalistic contexts for learning are visually complex and introduce challenges for young learners (Medina et al., 2011), but these real-world contexts may also provide opportunities to detect previously ignored regularities in the input. Existing research explains infants' ability to learn nouns, in particular, by exploiting consistent referent-linked visual cues to word meaning. However, this narrow conception of word knowledge as the ability to link labels with their

concrete, imageable, consistent visual referents cannot fully account for infants' early learning of everyday words, for which other cues may better predict learning.

Our video corpus analysis identified situational stability as one such potential cue. Even if *uh-oh* does not reliably co-occur with the visual referent of a fallen cup, for example, this word may appear consistently when a variety of objects fall, perhaps providing enough information for infants to detect this broader meaning-relevant pattern. The skewed nature of the situational-level input could support learning by facilitating efficient recognition of words that appear in their dominant situational contexts, while also promoting successful generalization to novel situations, given that meanings are not tightly constrained to a single context. Not unique to this set of everyday words, skewed distributions are a common feature of children's real-world linguistic (Goldberg et al., 2004; Montag et al., 2018) and visual input (Clerkin et al., 2017; Jayaraman et al., 2015) and have been identified as supportive of learning and generalization (Boyd & Goldberg, 2009; Carvalho et al., 2021; Maguire et al., 2008). Consistent with this view, we uncovered a preliminary link between regularity of the situational-level input and infants' rate of everyday word production (i.e., infants who heard a given target word in more consistent situational contexts may have been more likely to produce the word, relative to infants who saw the same word used in more variable situations).

It is worth noting, however, that while we aimed to avoid top-down assumptions, our categorization of situations for everyday words required adult decisions that may not map onto infants' parsing of complex, naturalistic scenes. Moreover, our decisions about categories may not map into infants' learning processes. This question of where to divide meaningful 'units' of information (i.e., between situational contexts that involve many dimensions of information, such as objects, actions, people, emotions, etc.) is not specific to the present study. For example, even

when operating within a single domain, like vision (Johnson et al., 2009; Slone & Johnson, 2015), speech (Frank et al., 2010; Mattys et al., 2005), or event processing (Baldwin & Kosie, 2021), there are similar questions surrounding the level at which learners segment and encode information. Still, it is promising that our annotated context-level features were able to capture potentially meaningful variation in infants' production (consistent with Roy et al., 2015). A key contribution of this work is the recognition that there may be learnable patterns in non-visual aspects of the input surrounding words, which remain largely overlooked in conventional studies of noun-to-object mapping. Future work is needed to determine how infants attend to and encode information in a given situation, and how this information is integrated across multiple episodes.

Additionally, it is worth noting that even within the set of early-learned everyday words investigated in the present study, there is considerable individual variation in the input. That is, while the general patterns of visual variability and situational stability appear for the five children and 12 words in our sample, idiosyncrasies at the child and item level cannot be ignored. Combined with evidence of idiosyncrasies in adult word representations (Wang & Bi, 2021), our findings suggest that early word meanings may be specific to infants' unique experiences. Word learning theories must account for individual trajectories (Samuelson, 2021) and acknowledge that word representations need only be similar enough from person to person to successfully communicate. Furthermore, theories must address how infants may be attending to a diverse array of cues in the environment, beyond object-specific visual characteristics. As identified in several previous studies, useful characteristics include overall label frequency (Goodman et al., 2008), frequency of occurrence in isolation (Brent & Siskind, 2001; Lew-Williams et al., 2011) or in sentence-final position (Braginsky et al., 2019), strong association with infancy (i.e., 'babiness,' Perry et al., 2015), short mean length of utterance and high contextual distinctiveness

(Roy et al., 2015), relatedness of word meanings (Floyd & Goldberg, 2021), and social cues such as eye gaze and referential gestures (Yu & Ballard, 2007). However, these features still cannot scale to a comprehensive theory of early word learning, and it remains unclear how multiple features are encoded and integrated within and across episodes. Together, recognized key predictors of word learning only account for up to 29% of the variance in age of acquisition across 386 English words on the MacArthur-Bates Communicative Development Inventory (CDI, Fenson et al., 1994) and notably even less for words that do not fit into established lexical categories (Braginsky et al., 2016).

Looking forward, continued rigorous manual coding efforts (Mendoza & Fausey, 2021) and advances in machine learning that allow for automatic processing of even larger amounts of naturalistic data (Orhan et al., 2020; Tsutsui et al., 2020) will position the field to be able to uncover novel features of infants' real-world input that may be supportive of word learning. Such work will greatly benefit from the increasing availability of modern video corpora that capture the first-person perspective of the learner across time (Bergelson et al., 2019; Sullivan et al., 2020). By embracing bottom-up approaches to describing the distributions of cues in the input, including multisensory, social, and contextual information, we can better understand how infants build early word representations.

References

- Baldwin, D. A., & Kosie, J. E. (2021). How does the mind render streaming experience as events? *Topics in Cognitive Science*, 13(1), 79–105.
- Bates, E., Marchman, V., Thal, D., Fenson, L., Dale, P., Reznick, J. S., Reilly, J., & Hartung, J. (1994). Developmental and stylistic variation in the composition of early vocabulary. *Journal of Child Language*, 21(1), 85–123.
- Bergelson, E., Amatuni, A., Dailey, S., Koorathota, S., & Tor, S. (2019). Day by day, hour by hour: Naturalistic language input to infants. *Developmental Science*, 22(1), e12715.
- Bergelson, E., & Swingley, D. (2012). At 6–9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences*, 109(9), 3253–3258.
- Bergelson, E., & Swingley, D. (2013). The acquisition of abstract words by young infants. *Cognition*, 127(3), 391–397.
- Bergelson, E., & Swingley, D. (2015). Early word comprehension in infants: Replication and extension. *Language Learning and Development*, 11(4), 369–380.
- Bloom, L., Tinker, E., & Margulis, C. (1993). The words children learn: Evidence against a noun bias in early vocabularies. *Cognitive Development*, 8(4), 431–450.
- Boersma, P., & Weenink, D. (2016). Praat software. *Amsterdam: University of Amsterdam*.
- Bowerman, M. (1978). The acquisition of word meaning: An investigation into some current conflicts. In *The development of communication* (pp. 263–287). Wiley.

- Boyd, J. K., & Goldberg, A. E. (2009). Input effects within a constructionist framework. *The Modern Language Journal*, 93(3), 418–429.
- Braginsky, M., Yurovsky, D., Marchman, V. A., & Frank, M. C. (2016). From *uh-oh* to *tomorrow*: Predicting age of acquisition for early words across languages. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.). *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 1691-1696). Cognitive Science Society.
- Braginsky, M., Yurovsky, D., Marchman, V. A., & Frank, M. C. (2019). Consistency and variability in children’s word learning across languages. *Open Mind*, 3, 52–67.
- Brent, M. R., & Siskind, J. M. (2001). The role of exposure to isolated words in early vocabulary development. *Cognition*, 81(2), B33–B44.
- Carvalho, P. F., Chen, C., & Yu, C. (2021). The distributional properties of exemplars affect category learning and generalization. *Scientific Reports*, 11(1), 1–10.
- Caselli, M. C., Bates, E., Casadio, P., Fenson, J., Fenson, L., Sanderl, L., & Weir, J. (1995). A cross-linguistic study of early lexical development. *Cognitive Development*, 10(2), 159–199.
- Clerkin, E. M., Hart, E., Rehg, J. M., Yu, C., & Smith, L. B. (2017). Real-world visual statistics and infants’ first-learned object names. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160055.
- Custode, S. A., & Tamis-LeMonda, C. (2020). Cracking the code: Social and contextual cues to language input in the home environment. *Infancy*, 25(6), 809–826.

- Çetinçelik, M., Rowland, C. F., & Snijders, T. M. (2021). Do the eyes have it? A systematic review on the role of eye gaze in infant language development. *Frontiers in Psychology, 11*, 3627.
- Demuth, K., Culbertson, J., & Alter, J. (2006). Word-minimality, epenthesis and coda licensing in the early acquisition of english. *Language and Speech, 49*(2), 137–173.
- Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D. J., Pethick, S. J., Tomasello, M., Mervis, C. B., & Stiles, J. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development, i*–185.
- Fernald, A., Perfors, A., & Marchman, V. A. (2006). Picking up speed in understanding: Speech processing efficiency and vocabulary growth across the 2nd year. *Developmental Psychology, 42*(1), 98.
- Fernald, A., Thorpe, K., & Marchman, V. A. (2010). Blue car, red car: Developing efficiency in online interpretation of adjective–noun phrases. *Cognitive Psychology, 60*(3), 190–217.
- Fernald, A., Zangl, R., Portillo, A. L., & Marchman, V. A. (2008). Looking while listening: Using eye movements to monitor spoken language. *Developmental Psycholinguistics: On-Line Methods in Children's Language Processing, 44*, 97.
- Floyd, S., & Goldberg, A. E. (2021). Children make use of relationships across meanings in word learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 47*(1), 29.
- Forbes, S. H., & Plunkett, K. (2019). Infants show early comprehension of basic color words. *Developmental Psychology, 55*(2), 240.

- Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2017). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, 44(3), 677–694.
- Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2021). *Variability and consistency in early language learning: The wordbank project*. MIT Press.
- Frank, M. C., Goldwater, S., Griffiths, T. L., & Tenenbaum, J. B. (2010). Modeling human performance in statistical word segmentation. *Cognition*, 117(2), 107–125.
- Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In S. Kuczaj II (Ed.) *Language Development* (Vol. 2): *Language, Thought and Culture*. Lawrence Erlbaum.
- Gleason, J. B., Perlmann, R. Y., & Greif, E. B. (1984). What's the magic word: Learning language through politeness routines. *Discourse Processes*, 7(4), 493–502.
- Gleason, J. B., & Weintraub, S. (1976). The acquisition of routines in child language. *Language in Society*, 5(2), 129–136.
- Goldberg, A. E., Casenhiser, D. M., & Sethuraman, N. (2004). Learning argument structure generalizations. *Cognitive Linguistics*, 15(3), 289–316.
- Golinkoff, R. M., Hirsh-Pasek, K., Cauley, K. M., & Gordon, L. (1987). The eyes have it: Lexical and syntactic comprehension in a new paradigm. *Journal of Child Language*, 14(1), 23–45.

- Goodman, J. C., Dale, P. S., & Li, P. (2008). Does frequency count? Parental input and the acquisition of vocabulary. *Journal of Child Language*, 35(3), 515–531.
- Jayaraman, S., Fausey, C. M., & Smith, L. B. (2015). The faces in infant-perspective scenes change over the first year of life. *PloS One*, 10(5), e0123780.
- Johnson, S. P., Fernandes, K. J., Frank, M. C., Kirkham, N., Marcus, G., Rabagliati, H., & Slemmer, J. A. (2009). Abstract rule learning for visual sequences in 8-and 11-month-olds. *Infancy*, 14(1), 2–18.
- Kartushina, N., & Mayor, J. (2019). Word knowledge in six-to nine-month-old norwegian infants? Not without additional frequency cues. *Royal Society Open Science*, 6(9), 180711.
- Kosie, J. E., & Baldwin, D. (2019). Attention rapidly reorganizes to naturally occurring structure in a novel activity sequence. *Cognition*, 182, 31–44.
- Lew-Williams, C., Pelucchi, B., & Saffran, J. R. (2011). Isolated words enhance statistical language learning in infancy. *Developmental Science*, 14(6), 1323–1329.
- Lieven, E. V., Pine, J. M., & Barnes, H. D. (1992). Individual differences in early vocabulary development: Redefining the referential-expressive distinction. *Journal of Child Language*, 19(2), 287–310.
- MacWhinney, B. (2000). *The CHILDES project: The database* (Vol. 2). Psychology Press.
- Maguire, M. J., Hirsh-Pasek, K., Golinkoff, R. M., & Brandone, A. C. (2008). Focusing on the relation: Fewer exemplars facilitate children’s initial verb learning and extension. *Developmental Science*, 11(4), 628–634.

- Mattys, S. L., White, L., & Melhorn, J. F. (2005). Integration of multiple speech segmentation cues: A hierarchical framework. *Journal of Experimental Psychology: General*, 134(4), 477.
- Medina, T. N., Snedeker, J., Trueswell, J. C., & Gleitman, L. R. (2011). How words can and cannot be learned by observation. *Proceedings of the National Academy of Sciences*, 108(22), 9014–9019.
- Mendoza, J. K., & Fausey, C. M. (2021). Quantifying everyday ecologies: Principles for manual annotation of many hours of infants’ lives. *Frontiers in Psychology*, 12(710636).
- Mirman, D. (2014). *Growth curve analysis and visualization using R*. Chapman; Hall/CRC.
- Montag, J. L., Jones, M. N., & Smith, L. B. (2018). Quantity and diversity: Simulating early word learning environments. *Cognitive Science*, 42, 375–412.
- Nastase, S. A., Goldstein, A., & Hasson, U. (2020). Keep it real: Rethinking the primacy of experimental control in cognitive neuroscience. *NeuroImage*, 222, 117254.
- Nelson, K. (1981). Individual differences in language development: Implications for development and language. *Developmental Psychology*, 17(2), 170.
- Orhan, A. E., Gupta, V. V., & Lake, B. M. (2020). Self-supervised learning through the eyes of a child. *Advances in Neural Information Processing Systems*, 33.
- Parise, E., & Csibra, G. (2012). Electrophysiological evidence for the understanding of maternal speech by 9-month-old infants. *Psychological Science*, 23(7), 728–733.

- Pereira, A. F., Smith, L. B., & Yu, C. (2014). A bottom-up view of toddler word learning. *Psychonomic Bulletin & Review*, 21(1), 178–185.
- Perry, L. K., Perlman, M., & Lupyan, G. (2015). Iconicity in english and spanish and its relation to lexical category and age of acquisition. *PloS One*, 10(9), e0137147.
- Ponari, M., Norbury, C. F., & Vigliocco, G. (2018). Acquisition of abstract concepts is influenced by emotional valence. *Developmental Science*, 21(2), e12549.
- Reuter, T., Dalawella, K., & Lew-Williams, C. (2021). Adults and children predict in complex and variable referential contexts. *Language, Cognition and Neuroscience*, 36(4), 474–490.
- Rose, Y., & MacWhinney, B. (2014). *The PhonBank project: Data and software-assisted methods for the study of phonology and phonological development*.
- Roy, B. C., Frank, M. C., DeCamp, P., Miller, M., & Roy, D. (2015). Predicting the birth of a spoken word. *Proceedings of the National Academy of Sciences*, 112(41), 12663–12668.
- Samuelson, L. K. (2021). Toward a precision science of word learning: Understanding individual vocabulary pathways. *Child Development Perspectives*, 15(2), 117–124.
- Slone, L. K., & Johnson, S. P. (2015). Infants' statistical learning: 2-and 5-month-olds' segmentation of continuous visual sequences. *Journal of Experimental Child Psychology*, 133, 47–56.
- Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, 106(3), 1558–1568.

- Stevens, J. S., Gleitman, L. R., Trueswell, J. C., & Yang, C. (2017). The pursuit of word meanings. *Cognitive Science*, 41, 638–676.
- Sullivan, J., Mei, M., Perfors, A., Wojcik, E., & Frank, M. C. (2020). SAYCam: A large, longitudinal audiovisual dataset recorded from the infant’s perspective. *Open Mind*, 1–10.
- Syrnyk, C., & Meints, K. (2017). Bye-bye mummy–word comprehension in 9-month-old infants. *British Journal of Developmental Psychology*, 35(2), 202–217.
- Tamis-LeMonda, C. S., Kuchirko, Y., Luo, R., Escobar, K., & Bornstein, M. H. (2017). Power in methods: Language to infants in structured and naturalistic contexts. *Developmental Science*, 20(6), e12456.
- Tardif, T., Fletcher, P., Liang, W., Zhang, Z., Kaciroti, N., & Marchman, V. A. (2008). Baby’s first 10 words. *Developmental Psychology*, 44(4), 929.
- Tincoff, R., & Jusczyk, P. W. (1999). Some beginnings of word comprehension in 6-month-olds. *Psychological Science*, 10(2), 172–175.
- Tincoff, R., & Jusczyk, P. W. (2012). Six-month-olds comprehend words that refer to parts of the body. *Infancy*, 17(4), 432–444.
- Tomasello, M. (1987). Learning to use prepositions: A case study. *Journal of Child Language*, 14(1), 79–98.
- Tsutsui, S., Chandrasekaran, A., Reza, M. A., Crandall, D., & Yu, C. (2020). A computational model of early word learning from the infant’s point of view. *arXiv*.

- Valleau, M. J., Konishi, H., Golinkoff, R. M., Hirsh-Pasek, K., & Arunachalam, S. (2018). An eye-tracking study of receptive verb knowledge in toddlers. *Journal of Speech, Language, and Hearing Research*, 61(12), 2917–2933.
- Wang, X., & Bi, Y. (2021). Idiosyncratic tower of babel: Individual differences in word-meaning representation increase as word abstractness increases. *Psychological Science*, 09567976211003877.
- Wojcik, E. H., Zettersten, M., & Benitez, V. L. (2022). The map trap: Why and how word learning research should move beyond mapping. *WIREs Cognitive Science*, e1596.
- Yu, C., & Ballard, D. H. (2007). A unified model of early word learning: Integrating statistical and social cues. *Neurocomputing*, 70(13-15), 2149–2165.

Supporting Online Information for “Moving beyond nouns in the lab”: Using naturalistic data to understand why infants’ first words include *uh-oh* and *hi*

The sections below provide supplemental analyses and figures, including:

- (1) Target word characteristics (Wordbank production percentiles and CHILDES frequency estimates)
- (2) Exploratory analyses of age, image familiarity, and item effects for Study 1 (Behavioral Experiment)
- (3) Visualization of the distribution of sampled tokens for Study 2 (Video Corpus Analysis), including across speakers and households
- (4) Figures depicting the frequency of corpus scene matches to experimental stimuli
- (5) Figures depicting the distribution of top visual referents, including variation across speakers and households
- (6) Figures depicting the distribution of top situational contexts, including variation across speakers and households
- (7) Examination of developmental change in the visual- and situational-level input surrounding everyday words

(8) Discussion of convergent findings from analyses of standard vs. variant forms of the target words

Analyses were conducted in R (version 4.1.0; R Core Team, 2021) using the *AICcmodavg* (version 2.3-1; Mazerolle, 2020), *broom* (version 0.7.9; Robinson et al., 2021), *childesr* (version 0.2.1.9000; Braginsky et al., 2021), *lme4* (version 1.1-27.1; Bates et al., 2015), *lmerTest* (version 3.1-3; Kuznetsova et al., 2017), *performance* (version 0.7.3; Lüdtke et al., 2021), and *tidyverse* (version 1.3.1; Wickham et al., 2019) packages, and plots were generated using *ggplot2* (version 3.3.5; Wickham, 2016), *grid* (version 4.1.0), and *gridExtra* (version 2.3; Auguie, 2017). **For access to all data, materials, and analysis scripts, see OSF (<https://osf.io/tdbqn/>).**

Study 1: Behavioral Experiment

Target word characteristics

Table S1 provides information about the 12 target everyday words. Production percentiles were collected from the Wordbank repository (Frank et al., 2017) and refer to the percentage of 16-month-old infants reported to understand and say the target word based on responses to the English (American) version of the Communicative Development Inventory: Words & Gestures (CDI; Fenson et al., 1994). Child Language Data Exchange System (CHILDES; MacWhinney, 2000) corpus frequency counts refer to the number of occurrences of the target words per 1 million words of speech input to North American English-hearing children, ranging in age from 0 to 17 months. Frequency counts include common variations of the target words (e.g., *bye-bye* includes *bye*, and *yum* includes *yummy*) and were extracted using the *childesr* package (Braginsky et al., 2021).

Table S1. Target word characteristics

Target word	CDI production percentile	CHILDES frequency
	(16 months)	(per 1 million words, up to 17 months)
<i>all-gone</i>	33	3,127
<i>bye-bye</i>	73	3,646
<i>hi</i>	67	7,227
<i>more</i>	32	7,975
<i>night-night</i>	42	562
<i>no</i>	63	21,656
<i>shh</i>	30	592
<i>thank-you</i>	39	4,777
<i>uh-oh</i>	75	9,908
<i>up</i>	35	15,111
<i>wow</i>	N/A	1,268
<i>yum</i>	59	1,101

GCA mixed-effects model details

We compared infants' LWL task performance against chance (0.5) by fitting a mixed-effects model predicting infants' accuracy in target looking at each time point in the critical window (367 to 4,000ms following target word onset) using the *lme4* package in R (Bates et al., 2015). The model included an intercept term and two orthogonal polynomial time terms for repeated measures (each 33-ms frame) as fixed effects as well as by-participant and by-item random intercepts. The dependent variable is the empirical log odds of the raw data, offset by 0.5 so that the intercept represents the proportion of target looking with chance at 0.5 (note that $\text{logit}(0.5) = 0$). Statistical significance for all model parameter estimates was assessed using the normal approximation (i.e., by treating the *t*-value as a *z*-value; Mirman, 2014).

Exploratory analyses of age and image familiarity effects

Following our pre-registered analysis plan, we explored whether age and image familiarity influenced the dynamics of looking behavior across the critical time window (367 to 4,000ms following target word onset). We fit a mixed-effects model, with the same base structure as above, predicting infants' accuracy in target looking at each time point in the critical window. Age (in days, scaled) and image congruence (strong/weak) were included as additional fixed effects, along with interactions between these predictors and all time terms. Image congruence was determined from parents' rating of how well the test images matched their child's typical visual experience when they heard the target words, reported on a scale from 1 [*completely unfamiliar*] to 4 [*very familiar*]. For each participant, we classified individual test images as either strongly congruent (3-4 on the scale) or weakly congruent (1-2 on the scale).

We first included by-participant and by-item random intercepts, as planned in our pre-registration, but when the initial model did not converge, we simplified the random effects structure following the recommendations outlined by Barr et al. (2013). Removing the by-item and by-participant random intercepts each led to a singular model fit. For transparency, the full model outputs are reported in Tables S2-S3.

Table S2. Fixed effects for growth curve model 1

Predictor	Estimate	SE	t
Intercept	0.52***	0.05	10.87
Linear time term	0.51**	0.18	2.89
Quadratic time term	0.02	0.16	0.10
Age (days, scaled)	0.02	0.04	0.48
Image congruence	-0.02	0.05	-0.36
Linear time term * Age	-0.37*	0.15	-2.48
Quadratic time term * Age	-0.02	0.18	-0.09
Linear time term * Image congruence	-0.87***	0.26	-3.31
Quadratic time term * Image congruence	-0.07	0.19	-0.36
Age * Image congruence	0.02	0.05	0.37
Linear time term * Age * Image congruence	0.26	0.29	0.88
Quadratic time term * Age * Image congruence	0.02	0.23	0.11

Note. The random effects structure for model 1 was as follows: (1 + Linear time term + Quadratic time term) * Age * Image congruence | Participant. Statistical significance for all model parameter estimates was assessed using the normal approximation (Mirman, 2014; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).

Table S3. Fixed effects for growth curve model 2

Predictor	Estimate	SE	t
Intercept	0.60***	0.08	7.88
Linear time term	0.23	0.16	1.24
Quadratic time term	0.34	0.18	1.87
Age (days, scaled)	0.02	0.05	0.31
Image congruence	-0.11*	0.05	-2.22
Linear time term * Age	-0.60**	0.16	-3.81
Quadratic time term * Age	-0.16	0.16	-1.01
Linear time term * Image congruence	-0.35	0.22	-1.61
Quadratic time term * Image congruence	-0.45*	0.22	-1.98
Age * Image congruence	0.03	0.04	0.77
Linear time term * Age * Image congruence	0.53**	0.18	2.89
Quadratic time term * Age * Image congruence	0.08	0.18	0.45

Note. The random effects structure for model 2 was as follows: (1 + Linear time term + Quadratic time term) * Age * Image congruence | Item. Statistical significance for all model parameter

estimates was assessed using the normal approximation (Mirman, 2014; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).

Exploratory analyses of item effects

To better understand the largely null results from the eye-tracking study, we also conducted exploratory analyses to look at item effects. We tested whether infants showed above-chance looking to target images individually for each of the 12 words using a one sample t -test and found that infants' accuracy in target looking was only significantly above chance for a single item: *hi* [$M = 0.74$, $t(32) = 5.37$, $p < 0.001$; Figure S1]. To control for potential bias associated with image preferences, we also computed item-pair mean difference scores (consistent with Bergelson & Swingley, 2013). This involved comparing infants' proportion looking to an image when it appeared as the target, relative to when it appeared as the unlabeled distracter, and thus controlled for any effects of differential salience between images in each yoked pair. Positive difference scores are indicative of word comprehension, and results showed that infants demonstrated above-chance levels of target looking for four items (*bye-bye*, *hi*, *shh*, *wow*). After correcting for multiple comparisons using the Holm-Bonferroni method, three items remained significantly above chance: *bye-bye* [$M = 0.21$, $t(32) = 3.51$, $p = 0.016$], *hi* [$M = 0.48$, $t(32) = 10.22$, $p < 0.001$], *wow* [$M = 0.16$, $t(32) = 3.18$, $p = 0.041$; Figure S2].

Figure S1. Mean accuracy by item (uncorrected) across all participants. Accuracy was calculated as the proportion of time that infants' fixated on the target image, relative to the total time that they spent looking at the target or distracter images in the critical window of analysis from 367 to 4000ms following target word onset. Points reflect means over participants. Chance-level performance (0.5) is indicated by the dashed line. Error bars reflect standard errors of the mean, and asterisks denote above-chance performance after correcting for multiple comparisons ($p < 0.004$).

Figure S2. Mean accuracy by item (corrected) across all participants. Accuracy was calculated as the proportion of time spent looking at a given test image when it appeared as the labeled target image, relative to the proportion of time spent looking at the same image when it appeared as the unlabeled distracter. Points reflect means over participants. Chance-level performance (0.0) is indicated by the dashed line. Error bars reflect standard errors of the mean, and asterisks denote above-chance performance after correcting for multiple comparisons ($p < 0.004$).

Exploratory analyses of vocabulary size effects

Infants' global receptive and expressive vocabulary knowledge was assessed using the CDI: Words & Gestures (Fenson et al., 1994), and survey responses were collected for 31 of 33 participants. Parents reported that, on average, infants understood 89.5 words ($SD = 84.1$, range = 7-396) and produced 8.2 words ($SD = 9.3$, range = 0-42). A correlational analysis found no relationship between infants' receptive ($r = 0.28$, 95% CI = [-0.08, 0.58], $p = 0.12$) or expressive ($r = 0.18$, 95% CI = [-0.18, 0.51], $p = 0.32$) vocabulary size and their accuracy in looking to the target images during the eye-tracking study.

Study 2: Video Corpus Analysis

Distributional statistics of sampled tokens

Table S4 and Figs. S3-S4 provide additional information about the 11,920 sampled tokens from the video corpus study, including the distribution of tokens across target words, speakers (child vs. adult), households, and individual recording sessions.

Table S4. Frequency of target words by child

Target word	Child name	Token count overall (child-produced)	Rate per session overall (child-produced)	Target word	Child name	Token count overall (child-produced)	Rate per session overall (child-produced)
<i>all-gone</i>	Alex	92 (28)	6.1 (1.9)	<i>shh</i>	Alex	44 (5)	2.9 (0.3)
	Lily	29 (1)	1.2 (0.04)		Lily	24 (9)	1.0 (0.4)
	Naima	178 (43)	4.3 (1.1)		Naima	35 (11)	0.9 (0.3)
	Violet	38 (1)	1.9 (0.1)		Violet	9 (0)	0.5 (0.0)
	William	120 (33)	8.6 (2.4)		William	48 (9)	3.4 (0.6)
	Average	91.4 (21.2)	4.4 (1.1)		Average	32.0 (6.8)	1.7 (0.3)
<i>bye-bye</i>	Alex	85 (24)	5.7 (1.6)	<i>thank-you</i>	Alex	203 (4)	13.5 (0.3)
	Lily	221 (37)	9.2 (1.5)		Lily	207 (7)	8.6 (0.3)
	Naima	212 (79)	5.2 (1.9)		Naima	232 (3)	5.7 (0.1)
	Violet	159 (54)	8.0 (2.7)		Violet	124 (4)	6.2 (0.2)
	William	85 (34)	6.1 (2.4)		William	82 (3)	5.9 (0.2)
	Average	152.4 (21.2)	6.8 (2.0)		Average	169.6 (4.2)	8.0 (0.2)
<i>hi</i>	Alex	201 (40)	13.4 (2.7)	<i>uh-oh</i>	Alex	280 (128)	18.7 (8.5)
	Lily	442 (27)	18.4 (1.1)		Lily	235 (51)	9.8 (2.1)
	Naima	371 (79)	9.1 (1.9)		Naima	237 (64)	5.8 (1.6)
	Violet	143 (21)	7.2 (1.1)		Violet	118 (17)	5.9 (0.9)
	William	183 (89)	13.1 (6.4)		William	191 (65)	13.6 (4.6)
	Average	268.0 (51.2)	12.2 (2.6)		Average	212.2 (65.0)	10.8 (3.5)
<i>more</i>	Alex	67 (1)	4.5 (0.1)	<i>up</i>	Alex	182 (16)	12.1 (1.1)
	Lily	106 (7)	4.4 (0.3)		Lily	263 (14)	11.0 (0.6)
	Naima	794 (278)	19.4 (6.8)		Naima	733 (180)	17.9 (4.4)
	Violet	73 (7)	3.7 (0.4)		Violet	223 (17)	11.2 (0.9)
	William	116 (12)	8.3 (0.9)		William	259 (47)	18.5 (3.4)
	Average	231.2 (61.0)	8.0 (1.7)		Average	332.0 (54.8)	14.1 (2.1)
<i>night-night</i>	Alex	21 (1)	1.4 (0.1)	<i>wow</i>	Alex	254 (76)	16.9 (5.1)
	Lily	56 (8)	2.3 (0.3)		Lily	282 (54)	11.8 (2.3)
	Naima	21 (9)	0.5 (0.2)		Naima	164 (4)	4.0 (0.1)
	Violet	18 (0)	0.9 (0.0)		Violet	52 (2)	2.6 (0.1)
	William	37 (7)	2.6 (0.5)		William	36 (10)	2.6 (0.7)
	Average	30.6 (5.0)	1.6 (0.2)		Average	157.6 (29.2)	7.6 (1.7)

<i>no</i>	Alex	632 (293)	42.1 (19.5)	<i>yum</i>	Alex	0 (0)	0.0 (0.0)
	Lily	664 (215)	26.7 (9.0)		Lily	95 (8)	4.0 (0.3)
	Naima	784 (365)	19.1 (8.9)		Naima	221 (25)	5.4 (0.6)
	Violet	297 (59)	14.9 (3.0)		Violet	7 (1)	0.4 (0.1)
	William	792 (379)	56.6 (27.1)		William	43 (21)	3.1 (1.5)
	Average	633.8 (262.2)	32.1 (13.5)		Average	73.2 (11.0)	2.6 (0.5)

Note. Token counts were summed across all recording sessions. Overall counts refer to tokens that were heard *or* produced by the target child. Child-produced counts refer to only those tokens produced by the target child. Rate per session was calculated by dividing the token count by the number of recording sessions, which varied across children (Alex-15, Lily-24, Naima-41, Violet-20, William-14)

Figure S3. Distribution of tokens across speakers.

Figure S4. Distribution of tokens across speakers, items, and sessions. Each point represents a single token, color coded by speaker. Sessions during which infants ranged in age from 11 to 24 months were coded. Because recordings began at the onset of first word production and families participated in a variable number of sessions, not all children have sessions that cover the entire 12-month timeframe.

Match to Study 1 stimuli

To evaluate the validity of our experimental stimuli, we calculated the proportion of corpus scene matches to the test images used in Study 1. The overall distribution of matches across all 12 items is shown in Figure S5. We also examined two potential sources of variability: differences between children's and adults' productions (Figure S6) and differences across households (Figure S7).

Figure S5. Overall frequency of corpus scene matches to Study 1 stimuli across items

Figure S6. Frequency of corpus scene matches to Study 1 stimuli across children's and adults' productions. To compare the distribution of matches to experimental stimuli for child vs. adult speakers (see Table S5 for summary statistics), we ran two-tailed, paired-samples *t*-tests but

found no significant differences across all three levels of match coding: visual matches [$t(11) = -0.85, p = 0.415$], situational matches [$t(11) = -0.31, p = 0.764$], and non-matches [$t(11) = 0.44, p = 0.668$].

Table S5. Summary of the frequency of corpus scene matches to Study 1 stimuli for children’s vs. adults’ productions

Match to stimuli	Child-produced tokens	Adult-produced tokens
	Mean prop. of matches (range)	Mean prop. of matches (range)
Visual match	9.2% (0-28.9%)	7.3% (0-18.1%)
Situational match	47.7% (1.5-84.8%)	46.6% (7.5-83.2%)
Non-match	43.1% (0-98.5%)	46.1% (5.5-87.6%)

Note. Match proportions reflect the average across items. Child-produced tokens refer to only those tokens produced by the target child, not other children (e.g., older siblings).

Figure S7. Frequency of corpus scene matches to Study 1 stimuli across children. Visual matches captured 7.0% of tokens, on average across words (range = 0-17.1%). Situational matches

captured 44.6% of tokens, on average (range = 8.9-84.0%). Non-matches captured 48.8% of tokens, on average (range = 5.1-85.6%).

Visual-level analyses

In our main analyses, we found that everyday words did not reliably co-occur with consistent visual referents (see Figure S8 for counts/proportions of unique visual referents). We also examined two potential sources of additional variability in co-occurring visual referents: differences between children's and adults' productions (Figure S9) and differences across households (Figure S10; Table S6).

Figure S8. (A) Raw counts of the number of unique visual referents for each item. Words co-occurred with 343 unique visual referents, on average. **(B)** Proportion of tokens corresponding to a unique visual referent for each item. Words appeared with a unique visual referent 34.5% of the time, on average.

Figure S9. Proportion of unique visual referents for children's vs. adults' productions. On average across items, 40.5% of child-produced tokens mapped onto a unique visual referent (range = 20.8-63.2%), relative to 41.2% for adult-produced tokens (range = 22.5-59.9%). To compare the proportion of unique referents for child vs. adult speakers, we ran two-tailed, paired-samples *t*-tests but found no significant difference in the level of visual variability across speakers [$t(11) = 0.20, p = 0.848$].

Figure S10. Distribution of top three visual referents across children. The top referent captured 7.4% of tokens, on average across words (range = 1.7-18.3%).

Table S6. Overlap in referents across children

Target word	Number of unique referents	Number of referents appearing for n children				
		5	4	3	2	1
<i>all-gone</i>	152	0	0	5	8	139
<i>bye-bye</i>	161	1	2	1	6	151
<i>hi</i>	247	5	4	4	20	214
<i>more</i>	314	0	4	7	25	278
<i>night-night</i>	68	0	0	0	5	63
<i>no</i>	1,414	1	1	10	37	1,365
<i>shh</i>	34	0	0	0	2	32
<i>thank-you</i>	357	1	3	3	20	330
<i>uh-oh</i>	344	2	1	7	21	313
<i>up</i>	659	1	0	7	28	622
<i>wow</i>	296	1	4	4	22	265
<i>yum</i>	74	N/A	1	1	6	66

Situational-level analyses

In our main analyses, we reported the overall distribution of top situational contexts () and suggested that the input surrounding everyday words may be skewed towards one or more dominant situations. Here, we also examined two potential sources of variability in the distribution of situational contexts: differences between children's and adults' productions (Figure S11) and differences across households (Figure S12).

Figure 11. Distribution of top three situational contexts for children's vs. adults' productions.

The top situation captured 50.0% of tokens, on average across words, for child-produced

utterances (range = 8.5-88.9%), and 46.8% for adult-produced utterances (range = 11.4-94.9%), with no significant differences based on age of the speaker [$t(11) = 0.56, p = 0.590$].

Figure S12. Distribution of top three situational contexts across children. The top situation captured 45.9% of tokens, on average across words (range = 12.8-96.6%).

Examining developmental change in the use of everyday words

Because children's language abilities improve dramatically over the sampled timeframe (e.g., Fenson et al., 1994; Fernald et al., 1998; Fernald et al., 2006) and past research suggests that child-directed input is modified to reflect such advances (e.g., Leung et al., 2021; Schwab et al., 2018; Yurovsky et al., 2016), we considered the possibility that the complexity of the visual and situational information associated with everyday words might change over time. Specifically, we asked whether everyday words became more visually variable and/or less situationally stable over time. **Visual variability** was calculated as the proportion of tokens corresponding to a unique visual referent for each item (e.g., for *uh-oh*, what proportion of the time did infants hear/use this word with a new visual referent?). **Situational stability** was measured as the proportion of tokens mapping onto the top situational context for each item (e.g., for *uh-oh*, what proportion of the time did infants hear/use this word to refer to any object falling?).

Our initial simple correlational analyses using age (in months) as the predictor variable found no significant increases in input complexity over time (Figure S13). To further explore the stability of these patterns, we re-ran the time-related analyses instead using time since the onset of expressive language, rather than age, to predict changes in the complexity of the input experienced by individual children. The age at which infants in the Providence corpus produced their first words ranged from 11 to 16 months ($M = 14$ months), suggesting individual differences in language ability that could confound analyses of developmental changes in input characteristics based solely on age, especially given the small sample size. Past research has found that, regardless of age, infants' experience producing language influences their ability to learn and recognize words (e.g., Fernald et al., 2001; Vihman et al., 2014; Werker et al., 2002), indicating that time since infants' first-word production could be a better predictor of

modifications to child-directed input in this dense corpus. In these updated analyses, we found the predicted pattern of increasing complexity of the input surrounding everyday words across the sampled timeframe (Figure S14). That is, everyday words appeared with more unique visual referents and the top situational context captured less of the input over time.

Still, these correlational analyses included only 12 data points (i.e., single point estimates for each time bin, averaged across words and children). To better account for potential variation across words and children, we fit linear mixed-effects models using the *lme4* package in R (Bates et al., 2015) to predict the variability of the visual-level and situational-level input surrounding everyday words as a function of time (in months) since the onset of expressive language. Models included by-child and by-item random intercepts, and statistical significance for the model parameter estimates was determined using the *lmerTest* package (Kuznetsova et al., 2017). We found that the complexity of the visual-level and situational-level input increased significantly across the sampled timeframe (Figure S15). That is, everyday words appeared with more unique visual referents ($Estimate = 0.01$, $SE = 0.003$, $t = 2.64$, $p = 0.009$) and the top situation captured slightly less of the input over time ($Estimate = -0.01$, $SE = 0.004$, $t = -4.00$, $p < 0.001$).

Figure S13. (A) Relationship between visual variability and age of the target child. The proportion of unique visual referents for each item was negatively correlated with age ($r = -0.64$, 95% CI = $[-0.88, -0.14]$, $p = 0.018$). (B) Relationship between situational stability and age of the target child. The proportion of tokens mapping onto the top situational context was not correlated with age ($r = 0.08$, 95% CI = $[-0.49, 0.60]$, $p = 0.795$).

Figure S14. (A) Relationship between visual variability and time since the onset of expressive language. The proportion of unique visual referents for each item increased linearly in the time following infants' first productions ($r = 0.69$, 95% CI = [0.22, 0.90], $p = 0.009$). (B) Relationship between situational stability and time since the onset of expressive language. The proportion of tokens mapping onto the top situational context decreased linearly in the time following infants' first productions ($r = -0.87$, 95% CI = [-0.96, -0.60], $p < 0.001$).

Figure S15. (A) Change in visual variability across time, operationalized as the proportion of tokens mapping onto a unique visual referent. (B) Change in situational stability across time, operationalized as the proportion of tokens mapping onto the top situational context. Individual points refer to by-child and by-item proportion estimates. Dot sizes correspond to the total number of tokens sampled for that child/item pair at each time point.

Comparing results for standard vs. variant form of target words

In our main analyses, we collapsed across standard and variant forms of each of the 12 target words (e.g., token counts for *bye-bye* included *bye* and *goodbye*; see Table S7 for full list of coded variants). For each item, 35.74% of tokens, on average, were considered a variant form (range = 0-72.13%; Figure S16). To ensure that the inclusion of variants did not artificially inflate estimates of visual or situational variability in our characterization of the input surrounding everyday words, we re-ran our descriptive analyses on just the subset of tokens matching the standard form of the target word. Results were qualitatively similar for the full set and the standard-only subset in terms of the number of unique visual referents per item (Figure S17), the relatively low proportion of tokens captured by the top visual referents (Figure S18), the greater proportion of tokens captured by the top situational contexts (Figure S19), and the frequency of corpus scene matches to our experimental stimuli (Figure S20).

Figure S16. Frequency of standard vs. variant form of target words.

Table S7. List of coded variants of target words

Target word	Coded forms	Target word	Coded forms
<i>all-gone</i>	all-gone, all-done, gone, done	<i>shh</i>	shh, shush, hush
<i>bye-bye</i>	bye-bye, bye, goodbye	<i>thank-you</i>	thank-you, thanks
<i>hi</i>	hi, hello	<i>uh-oh</i>	uh-oh, oh-no, oops(ie), whoops(ie)
<i>more</i>	more	<i>up</i>	up, uppa
<i>night-night</i>	night-night, goodnight, 'night	<i>wow</i>	wow, whoa
<i>no</i>	no, nope	<i>yum</i>	yum, yummy

Figure S17. (A) Raw counts of the number of unique visual referents for each item (standard form only). Words co-occurred with 288 unique visual referents, on average. (B) Proportion of tokens corresponding to a unique visual referent for each item (standard form only). Words appeared with a unique visual referent 37.1% of the time, on average.

Figure S18. Distribution of top visual referents across target words (standard form only). The top referent captured 12.0% of tokens, on average (range = 1.5-25.5%).

Figure S19. Distribution of top situational contexts across target words (standard form only). The top situation captured 48.4% of tokens, on average (range = 12.8-94.1%). As with the combined standard/variant analyses, the top three situational contexts captured a majority of tokens for 9/12 items, with proportions ranging from 31.2-100% of tokens ($M = 73.8\%$).

Figure S20. Frequency of corpus scene matches to Study 1 stimuli across items (standard form only). Visual matches captured 7.6% of tokens, on average across words (range = 0-20.7%). Situational matches captured 46.9% of tokens, on average (range = 7.3-80.7%). Non-matches captured 45.5% of tokens, on average (range = 5.9-86.5%).

References

- Auguie, B. (2017). *gridExtra*: Miscellaneous functions for “grid” graphics. R package version 2.3. <https://CRAN.R-project.org/package=gridExtra>
- Barr, D. J., Levy, R., Scheepers, C., & Tilly, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255-278. <https://doi.org/10.1016/j.jml.2012.11.001>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Bergelson, E., & Swingley, D. (2013). The acquisition of abstract words by young infants. *Cognition*, 127(3), 391–397. <https://doi.org/10.1016/j.cognition.2013.02.011>
- Braginsky, M., Sanchez, A., & Yurovsky, D. (2021). *childesr*: Accessing the ‘CHILDES’ database. R package version 0.2.1.9000. <https://github.com/langcog/childesr>
- Fenson, L., Dale, P. S., Reznick, J. S., Bates, E., Thal, D. J., Pethick, S. J., Tomasello, M., Mervis, C. B., & Stiles, J. (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development*, 59(5), i–185. <https://doi.org/10.2307/1166093>
- Fernald, A., Perfors, A., & Marchman, V. A. (2006). Picking up speed in understanding: Speech processing efficiency and vocabulary growth across the 2nd year. *Developmental Psychology*, 42(1), 98-116. <https://doi.org/10.1037/0012-1649.42.1.98>

- Fernald, A., Pinto, J. P., Swingley, D., Weinberg, A., & McRoberts, G. W. (1998). Rapid gains in speed of verbal processing by infants in the 2nd year. *Psychological Science*, 9(3), 228-231. <https://doi.org/10.1111/1467-9280.00044>
- Fernald, A., Swingley, D., & Pinto, J. P. (2001). When half a word is enough: Infants can recognize spoken words using partial phonetic information. *Child Development*, 72(4), 1003-1015. <https://doi.org/10.1111/1467-8624.00331>
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). *lmerTest* package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1–26. <https://10.18637/jss.v082.i13>
- Leung, A., Tunkel, A., & Yurovsky, D. (2021). Parents fine-tune their speech to their children's vocabulary knowledge. *Psychological Science*. <https://doi.org/10.1177/0956797621993104>
- Lüdtke, D., Ben-Shachar, M. S., Patil, I., Waggoner, P., & Makowski, D. (2021). *performance*: An R package for assessment, comparison, and testing of statistical models. *Journal of Open Source Software*, 6(60), 3139. <https://doi.org/10.21105/joss.03139>
- Mazerolle, M. J. (2020). *AICcmodavg*: Model selection and multimodel inference based on (Q)AIC(c). R package version 2.3-1. <https://cran.r-project.org/package=AICcmodavg>
- R Core Team (2021). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org>
- Robinson, D., Hayes, A., & Couch, S. (2021). *broom*: Convert statistical objects into tidy tibbles. R package version 0.7.9 <https://CRAN.R-project.org/package=broom>

- Schwab, J. F., Rowe, M. L., Cabrera, N., & Lew-Williams, C. (2018). Fathers' repetition of words is coupled with their children's vocabularies. *Journal of Experimental Child Psychology*, 166, 437-450. <https://doi.org/10.1016/j.jecp.2017.09.012>
- Vihman, M. M., DePaolis, R. A., & Keren-Portnoy, T. (2014). The role of production in infant word learning. *Language Learning: A Journal of Research in Language Studies*, 64(s2), 121-140. <https://doi.org/10.1111/lang.12058>
- Werker, J. F., Fennell, C. T., Corcoran, K. M., & Stager, C. L. (2002). Infants' ability to learn phonetically similar words: Effects of age and vocabulary size. *Infancy*, 3(1), 1-30. https://doi.org/10.1207/S15327078IN0301_1
- Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. New York: Springer-Verlag.
- Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L. D., François, R., ... Yutani, H. (2019). Welcome to the tidyverse. *Journal of Open Source Software*, 4(43), 1686. <https://doi.org/10.21105/joss.01686>
- Yurovsky, D., Doyle, G., & Frank, M. C. (2016). Linguistic input is tuned to children's developmental level. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.). *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 2093-2098). Austin, TX: Cognitive Science Society.