ATTENTION TO SHARED PERCEPTUAL FEATURES INFLUENCES EARLY NOUN-CONCEPT PROCESSING

by

Ryan E. Peters

A Dissertation

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy

Department of Speech, Language, & Hearing Sciences

West Lafayette, Indiana

August 2019

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL

Dr. Arielle Borovsky, Chair

Department of Speech, Language, and Hearing Sciences

Dr. Laurence Leonard

Department of Speech, Language, and Hearing Sciences

Dr. Amanda Seidl

Department of Speech, Language, and Hearing Sciences

Dr. Brandon Keehn

Departments of Speech, Language, and Hearing Sciences and Psychological Sciences

Approved by:

Dr. Keith Kluender

Head of the Graduate Program

For Fumi, who helped me follow my dream halfway around the world.

ACKNOWLEDGMENTS

This dissertation would not exist if not for the great deal of assistance and encouragement I received over the past five years. This work was supported by the NIDCD under award numbers F31 DC017089, awarded to the author, and R03 DC013638, awarded to Dr. Arielle Borovsky. Beyond funding, I am grateful to Dr. Borovsky for being an incredible mentor and collaborator. She provided me with more opportunities than I could have ever hoped for, and her guidance has made this a remarkably rewarding journey. I would also like to thank my dissertation committee members at both Florida State University and Purdue University – Drs. Don Compton, Michael Kaschak, Brandon Keehn, Colleen Kelley, Laurence Leonard, and Amanda Seidl – for their support as I formed an idea and then turned it into a completed study. In addition, I would like to thank Dr. Ken McRae for his guidance in creating the feature norm data set that made this work possible. Next, Justin Kueser has been an amazing colleague and friend, who spent countless hours listening to me talk about my research and providing invaluable feedback. I would also like to thank the members of the LCD, Purdue Infant Speech, and LLaMA labs – in particular Elizabeth Allen, Emanuel Boutzoukas, Nicky Hopkins, Michelle Indarjit, Kenzie McAfee, and Brooke Reffeitt – for their invaluable help in creating and running my studies. Finally, I am grateful for my friends and family – especially my parents, my sister Stacy, my wife Fumiko, my son Akala, and our scruffy dog Moja – who were always there for me when I most needed support, encouragement, and some revitalizing time outdoors.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Author: Peters, Ryan, E. PhD
Institution: Purdue University
Degree Received: August 2019
Title: Patterns of Shared Perceptual Features Influence Early Noun-Concept Processing
Committee Chair: Arielle Borovsky

Recent modeling work shows that patterns of shared perceptual features relate to the group-level order of acquisition of early-learned words (Peters & Borovsky, 2019). Here we present results for two eye-tracked word recognition studies showing patterns of shared perceptual features likewise influence processing of known and novel noun-concepts in individual 24- to 30-month-old toddlers. In the first study (Chapter 2, N=54), we explored the influence of perceptual connectivity on both initial attentional biases to known objects and subsequent label processing. In the second study (Chapter 3, N=49), we investigated whether perceptual connectivity influences patterns of attention during learning opportunities for novel object-features and object-labels, subsequent pre-labeling attentional biases, and object-label learning outcomes. Results across studies revealed four main findings. First, patterns of shared (visual-motion and visual-form and surface) perceptual features do relate to differences in early noun-concept processing at the individual level. Second, such influences are tentatively at play from the outset of novel noun-concept learning. Third, connectivity driven attentional biases to both recently learned and well-known objects follow a similar timecourse and show similar patterns of individual differences. Fourth, initial, pre-labeling attentional biases to objects relate to subsequent label processing, but do not linearly explain effects of connectivity. Finally, we consider whether these findings provide support for shared-feature-guided selective attention to object features as a mechanism underlying early lexico-semantic development.

CHAPTER 1. INTRODUCTION

Individual differences in early language abilities persist over time and robustly predict a range of long-term outcomes. For example, low toddler vocabulary size predicts lower performance in a wide range of language-related skills in childhood and adolescence (Rescorla, 2002, 2009), and increased risk of developing clinically significant language and learning impairments, such as dyslexia (Lyytinen, Eklund, & Lyytinen, 2005) and Developmental Language Disorder (Dale, Price, Bishop, & Plomin, 2003). Developmental Language Disorder is, in turn, a risk factor for compounding negative outcomes ranging from poor academic achievement and professional attainment to impaired social and emotional functioning (Bishop & Leonard, 2000; Bronwlie et al., 2004; Conti-Ramsden & Durkin, 2016). While these predictive links are well documented, the mechanisms underlying individual differences in early vocabulary size remain poorly understood and are an important line of ongoing research. Recent findings that separately tie *patterns of semantic structure* in individual children's early productive vocabularies to *word processing skill* on the one hand (e.g., Borovsky, Ellis, Evans, & Elman, 2015; 2016), and use of *attentional biases* to perceptual features of objects on the other (e.g., Colunga & Sims, 2017), hint at an interesting possibility. Namely, noun-concept processing may relate to the semantic connectivity of a toddler's vocabulary via attentional biases that depend on relations between the semantic features of to-be-processed words and known vocabulary. If this turns out to be the case, then targeted interventions based on individual children's vocabulary structures may be a way to boost long-term outcomes of at-risk children with smaller vocabularies.

In this dissertation, we explore how the structure of shared perceptual features interconnecting the growing lexico-semantic networks in toddlers' minds influence early noun-concept processing. The structure will be as follows. In this introductory chapter, we first discuss research on relations between toddlers' word meaning knowledge, word processing skill, and attentional biases. We then introduce a useful methodology: modeling early lexico-semantic knowledge using noun-feature networks. Next, we review research using this methodology to explore early word processing and learning, concluding with a recent project demonstrating perceptual features matter most in the mind of a toddler. Chapters 2 and 3 pick up this thread and present empirical explorations of how patterns of shared perceptual features influence toddlers' processing of known and novel noun-concepts. Finally, in the concluding chapter, we place the findings within the broader picture laid out in the introduction, and plot the way forward.

Word Meaning Knowledge, Processing Skill, and Attentional Biases

In this section we lay the groundwork for the hypothesis that noun-concept processing skill relates to the semantic connectivity of a toddler's vocabulary via attentional biases that depend on relations between the semantic features of to-be-processed words and known vocabulary. We begin by covering work exploring relations between the structure of toddlers' word meaning knowledge and their word processing skill. We next introduce work relating vocabulary composition to the use of attentional word learning biases and then discuss how similar mechanisms could underly individual differences in word processing skill more broadly.

Individual differences in vocabulary composition and word processing skill

Children's early word learning environments contain structure across a variety of dimensions, certain aspects of which they leverage to learn language. For some dimensions, such as those including word sounds and word meanings, an important part of the leveraging process is the incorporation of the structure into the representations within the minds of the learners. For example, for word meanings, children are able to recognize patterns of similarity relations between objects and form categories (e.g., Cohen & Cashon, 2006; Madole & Oakes, 1999). Furthermore, children are able to recognize new category members and leverage their category knowledge to learn new words. Borovsky and colleagues (2015) explored this process in 2-year-old toddlers. First, they measured each toddler's word knowledge using the MacArthur-Bates Communicative Developmental Inventory (MBCDI; Fenson et al., 2006), a commonly used parental questionnaire of productive vocabulary. Then, for each child, they split 6 categories of objects (animals, body parts, clothing, drinks, fruits, vehicles) into 3 high-density and 3 low-density categories based on the proportions of words known in each of the categories. Finally, they measured word processing using the 'looking while listening' method (Fernald, et al., 2008), in which eye-gaze patterns to two objects are tracked while the participant hears one of them labeled. The results indicated that high category density facilitated the learning of novel items. Interestingly, this facilitatory effect of high-category density held true for known items as well (Borovsky, Ellis, Evans, & Elman, 2016). In other words, within individuals, the structure of word meanings incorporated into their own lexico-semantic systems differentially affect processing of both known and novel noun-concepts.

How might membership in higher-density categories facilitate processing? One explanation builds on work exploring how words' phonological *neighborhoods* influence learning. Words that share (either phonological or semantic) features are often described as *neighbors*. Words with a large number of phonological neighbors, or in other words have a high phonological neighborhood density, are learned more easily by adults (Storkel, Armbruster &

Hogan, 2006; Storkel, Bontempo & Pak, 2014), preschoolers (Storkel, 2001), and toddlers (Newman, Samuelson, & Gupta, 2008). Though for children there are conflicting results (e.g., Gray, Pittman, & Weinhold, 2014; Swingley & Aslin, 2007), perhaps related to individual differences in vocabulary size (Storkel & Hoover, 2011). The facilitatory effect of high phonological neighborhood density is thought to arise when participants recognize features in the input that are shared with known items, which in turn aids the integration of new lexical representations with existing representations (Storkel et al., 2006). Could something similar happen for words in high-density semantic neighborhoods? Neural network simulations (Borovsky & Elman, 2006) duplicating the aforementioned facilitatory effect of high category density support this hypothesis: networks with more interconnected semantic structure were more easily able to recognize and map similarities from known to novel items, and in turn more easily integrated novel items into the networks' internal representations. Furthermore, the facilitated learning in highly structured networks compounded over time, resulting in large differences in vocabulary size compared to less-structured networks. Notably, the initial step in the process – recognizing features shared with known words – approximately maps onto a well-studied mechanism in early word learning: the selective attention to relevant object features when generalizing names to new objects.

Selective attention to relevant features in early word learning

The featural structure of objects in different domains can vary greatly. For example, for count nouns such as SPOON, object shape is a defining feature, while material is less important (i.e., spoons can be metal, wood, or plastic). In contrast, for mass nouns such as ICE, material is the defining feature, while shape is irrelevant. In the novel-noun generalization task – where participants are presented with a novel object, told its name (e.g., "This is a 'dax'."), and then asked which of two additional objects (that are related to the first in different ways) are also called a dax – adults and older children reliably generalize along the shape dimension for count nouns, a tendency which has been described as the *shape bias* (Landau, Smith, & Jones, 1988). Selective attention to shape depends on patterns in the input, such as task or linguistic context and the featural structure of objects. For example, biased attention to shape is robustly evident when children are learning object names but *not* when judging category membership or degree of similarity between objects (Landau, Smith, & Jones, 1988; Imai, Gentner, & Uchida, 1994; Soja, Carey, & Spelke, 1991). In other words, "young children do not go around in their everyday lives always attending only to the shape of objects. Rather, the shape bias appears to reflect specifically a child's knowledge about how words map onto object categories." (p. 54, Smith, 2000). Furthermore, when other features are present that are highly relevant to a category, children shift their attentional biases in line with what is relevant to that category. For example, when eyes or feet are added to the same objects for which children show a robust shape bias – signifying the objects are now animals – 3-year-olds only generalize names to objects which share *both* shape and material (Jones, Smith, & Landau, 1991; Jones & Smith, 1998). Based on these facts, Smith and colleagues (e.g., Smith & Colunga, 2008; Smith & Samuelson, 2006) have proposed an *Attentional Learning Account* to explain the patterns of generalizations and biases that children exhibit in early word learning. A key point is that shape is not important in and of itself, but instead becomes important in naming contexts for certain domains of objects because of the co-occurrence of naming, object features, and the relational structuring of objects within categories or domains across an individual's past word learning and processing episodes. It naturally follows that such attentional word learning biases develop over the course of early word learning as children gain experience with relevant patterns in the input.

Individual differences in vocabulary composition and use of attentional biases

While there does not appear to be a specific age at which individuals reliably begin to use the shape bias, nor a "critical mass" of count nouns that must be known (Gershkoff-Stowe & Smith, 2004), there is a normative developmental trajectory from lack of selective attention when generalizing names, through highly variable usage of a shape bias, to systematic usage of a shape bias. 13-month-olds show no shape bias, and instead only reliably generalize newly learned names to objects that are identical (Woodward, Markman, & Fitzsimmons, 1994). Slightly older children extend novel labels to objects that are different but share perceptual similarities – such as being a different color (Woodward et al., 1994) or being a member of the same category (e.g., vegetables; Waxman & Hall, 1993). Selective attention to shape develops around 2-years-of-age in typically developing children – when children have between 50 and 150 nouns in their productive vocabulary (Gershkoff-Stowe & Smith, 2004). However, usage of the shape bias is initially highly variable and continues to increase in certainty and selectivity across the third year of life (e.g., Landau et al., 1988, 1992, 1998; Imai & Gentner, 1997, Imai, Gentner, & Uchida, 1994).

Individual differences in usage of the shape bias are most apparent just as such attentional word learning biases begin to come online, in the latter half of the second year of life (Colunga & Sims, 2017; Jones, 2003; Perry & Samuelson, 2011). This is particularly true for children at the lower end of the vocabulary size spectrum, who have vocabularies with more variable composition than children with larger vocabularies. For example, Perry and Samuelson (2011) found a robust relation between vocabulary composition and usage of the shape bias, but only for children who knew less than 50 nouns. Likewise, Colunga and Sims (2017) presented computational and empirical evidence that neural network models and children with high vocabularies show a normative pattern of feature biases (i.e., they show a shape bias for solids and a material bias for non-solids), but models and children with low vocabularies show more variable patterns – with some showing the normative patterns, while others show mixed patterns such as overgeneralization of the shape bias to non-solids.

While early differences in word learning appear to influence the development of attentional word learning biases, the reverse is also true: teaching infants to pay attention to shape (or other properties such as material) influences both the type of words children learn and their rate of word learning (Colunga & Sims, 2011; Smith, Jones, Landau, Gershkoff-Stowe, & Samuelson, 2002). Importantly, such differences in the rate of word learning due to differences in the usage of selective attentional biases could compound over time, contributing to long-term outcomes. Indeed, support for this possibility comes from work showing that children with Developmental Language Disorder show very different patterns of attentional word learning biases from their typically developing peers (Collison et al., 2015). Altogether, it is clear that selective attentional word learning biases have bi-directional relations with the vocabulary composition of learners, and emerge over time through experience (Smith, 2005; Smith, Jones, & Landau, 1996; Smith & Samuelson, 2006; Yoshida & Smith, 2003). This aligns well with the viewpoint that children are active learners with a variety of attentional biases based on their current knowledge that mediate their online interactions with, and thus what they learn from, the environment (Bates & MacWhinney, 1987; MacWhinney, 1987; Smith, 2000). Furthermore, this aligns with the neural network modeling work that duplicated the facilitatory effect of high category density (Borovsky & Elman, 2006), in that the effect both emerged from the pattern of known words and had subsequent downstream effects on vocabulary size outcomes. In the next section we describe how these effects might arise due to context guided selective attention similar to the shape bias.

Might feature biases influence early word processing more generally?

The Attentional Learning Account (ALA; e.g., Colunga & Smith, 2008; Smith & Samuelson, 2006) is specifically designed to explain patterns of name generalization to novel objects, but it is built upon two of the most fundamental mechanisms in psychology: associative learning and the automatic control of selective attention by predictive cues. As discussed above, the central premise of the ALA is that children recognize coherent structure across linguistic and nonlinguistic input in naming situations, and then (via associative learning) heuristically capitalize on this structure when learning words and selectively attend to object features that have correlated with similar linguistic contexts in their experience. Notably, the shape bias is not necessarily the first or simplest selective attentional bias, but rather one that is important because it is broadly beneficial for learning a wide array of noun-concepts. Indeed, as demonstrated by the aforementioned work showing differential use of the shape bias when objects have eyes or feet (Jones et al., 1991; Jones & Smith, 1998), children employ a variety attentional biases that are specific to different categories of objects, with shifts in attention triggered by the recognition of salient features that have proven predictive across past experiences.

Relevant to the main questions at hand, this work can be logically extended to explain processing differences that relate to knowing more or less words from a given category. For example, take a child who has eaten a variety of objects, learned their names, and learned to distinguish them visually by color and texture via the recognition of coherent covariation (Rogers & McClelland, 2004) in these properties across past experiences with edible objects. Upon noticing that a person is eating an unfamiliar object, that child might shift her attention to the object’s color and texture. In other words, the recognition of a shared feature could act as a predictive cue that directs attention to other featural dimensions that have associated across past experiences. In turn, by systematically attending to potentially relevant featural dimensions of the object she has the potential to jump start the process of building a cohesive representation with structure influenced by other known edible items (Jones et al., 1991). Such correlation structure within a representation means that any subsequently added information will be automatically enmeshed within a web of mutually predictive relations (Smith, Colunga, & Yoshida, 2010). Thus, if someone happens to label the aforementioned novel edible object, the child’s object-label mapping could be boosted relative to another child with less experience with foods and who did not employ feature-guided selective attention. To reiterate, children who know a variety of items in a given category might be better able to employ selective attention based on the structure of features across items in the category when encountering potential new category members. Such feature-guided selective attention may jump start the creation of a cohesive representation, which in turn could be beneficial for subsequent object-label mapping. Importantly, the benefits of such processes act over multiple timescales, and could extend beyond the initial name learning episode, snowballing across subsequent naming interactions with objects and resulting in even richer, more jointly interactive sets of mutually predictive featural cues that result in more easily accessible representations (Colunga & Smith, 2008; Smith et al., 2010). Furthermore, considering word learning is an ongoing process that continues well past the first exposure and relates to improvements in familiar word recognition (McMurray, Horst, & Samuelson, 2012), this explanation cleanly ties together the facilitatory effect of high category density for both novel and known object recognition. Specifically, *novel* objects in high-density categories may benefit from an initial feature-guided selective attention derived processing boost, while *known* objects in high-density categories may have a history of such facilitated processing – which in turn has pushed the words further along the familiarity continuum then known words in low-density categories, resulting in relatively facilitated processing.

This account leaves one thing somewhat unclear: the timing. In both the empirical evidence and the proposed mechanism, there are two time points where the facilitatory effect might take place. For example, in regard to feature-guided selective attention, the recognition of the initial feature (eating in the previous example) might not only serve to guide attention and structure subsequent representations, but might also bring attention to the object in the first place. Likewise, for the facilitatory effect of high category density, it is unclear if the effect is driven by biased attention to the object pre-labeling, or facilitated processing of the label. For now, we will set this issue aside, but we will return to it in the empirical studies described in chapters 2 and 3.

The ideas presented in this section build the groundwork for the primary hypothesis of this dissertation. Namely, noun-concept processing skill relates to the composition of a toddler's vocabulary via selective attentional biases driven by features shared between to-be-processed words and known vocabulary. The foundation is built upon work relating the structure of toddlers' word meaning knowledge and their word processing skill on the one hand, and work relating vocabulary composition to the use of attentional word learning biases on the other. Finally, these bodies of work are tied together by an application of the Attentional Learning Account (Smith & Samuelson, 2006) to describe shared-feature-guided selective attention to explain facilitatory effects of high-category density on novel and known word processing. A key point of this extension is that, rather than a broadly applicable mechanism along the lines of the shape bias, such shared-feature-guided processes will be more idiosyncratic to clusters of known objects that form (potentially ad-hoc) categories. Furthermore, given the category density effect, it is entirely possible, or perhaps even likely, that individual toddlers use a range of such processes for each cluster of objects they have recognized within their collection of known objects. In other words, to explore this issue we need to move beyond simple descriptions of categories or features, and work with more granular featural descriptions of objects to explore patterns of shared features across the full breadth of individual toddlers' vocabularies. Therefore, in the following sections we introduce the methodology of modeling toddlers' vocabularies using noun-feature networks – an approach that builds on recent advances in the application of graph-theoretic network modeling techniques to lexico-semantic representations (e.g., Hills et al., 2009a, 2009b; Steyvers, & Tenenbaum, 2005).

Characterizing Lexico-Semantic Structure Using Noun-Feature Networks

In this section we introduce feature-based lexico-semantic networks as a useful tool for characterizing both the structure of toddlers' productive vocabularies and the granular featural makeup of individual words. We start by tracing the roots of the modern conceptualization of feature-based semantic networks, before then introducing the central concepts necessary for modeling and characterizing such networks.

Historical foundations

The modern conceptualization of word meaning knowledge as a lexico-semantic network emerged from work on semantic memory organization (Collins & Quillian, 1969; Quillian, 1967, 1969). Collins and Loftus (1975) hypothesized that semantic cognition could be explained by *spreading activation* that flows along the semantic connections between concepts. This hypothesis was supported by semantic priming experiments, in which the processing of a target word is speeded if presented after a related word (e.g., Neely, 1977, 1991). Initially, the structure of semantic networks was hypothesized to be of a superordinate-subordinate hierarchical nature, with lower level concepts inheriting all the properties of their respective superordinate concepts. However, evidence of the graded nature of superordinate category membership and the privileged status of some concepts revealed relations between concepts are far more complicated and nuanced (Rips, Shoben, & Smith, 1973; Rosch & Mervis, 1975; Rosch, Simpson, & Miller, 1976). Out of this research emerged the idea that the connections between concepts – and the resulting large-scale structure – are largely similarity based, with one dominant definition of similarity being the degree of overlap in feature-based semantic representations of meanings.

Researchers classify the subcomponents of word (e.g., SPOON) meanings, using semantic features (e.g., *has a handle, used for eating, a utensil*, etc.) that are either created by the researcher (e.g., Rumelhart, 1990) or empirically derived by collecting semantic feature production norms (e.g., McRae, Cree, Seidenberg, & McNorgan, 2005). Feature production norms are collected by presenting participants with a list of concepts and asking them to provide properties that come to mind for each concept. Participants' responses are often idiosyncratic, so words are given to many participants and only those that appear repeatedly make it into the final feature production norm dataset. This procedure is argued to provide insight into conceptual representations "not because they yield a literal record of semantic representation (i.e., they do not provide a verbatim read-out), but rather because such representations are used systematically by participants when generating features (Barasalou, 2003)" (McRae et al., 2005, p. 549). Simply put, feature norms are not definitive psychological representations, but they have been used to explain a variety of findings in adult psycholinguistic processing, including semantic richness effects (Pexman, Hargreaves, Sialuk, Bodner, & Pope, 2008; Rabovsky, Sommer, & Abdel

Rahman, 2012), conceptual structure (Hills, Maouene, Maouene, Sheya, & Smith, 2009b; Medin, 1989), semantic priming (McRae, De Sa, & Seidenberg, 1997; Vigliocco, Vinson, Lewis, & Garret, 2004), categorization (Hampton, 1997; Smith, Shoben, & Rips, 1974), and word/concept learning (Rogers & McClelland, 2004). Of key relevance to the issues laid out in the previous section, such feature norms can be used to create graph-theoretic noun-feature networks. In the next section, I introduce this technique and review research that uses it to explore early lexico-semantic development.

Central concepts of graph-theoretic noun-feature network modeling

The foundation of graph-theoretic network modeling rests on the definition of a network, otherwise known as a *graph*, as a set of *nodes* (aka *vertices* or *points*) that are connected via a set of pairwise *links* (Figure 1). The reduction of complex systems to a set of nodes and links makes otherwise hopelessly complicated issues tractable, and allows for comparisons between things as different as vocabulary structure and power grids (Barabasi, 2015).

In typical graph-theoretic representations of lexico-semantic networks, nodes consist of words and links consist of semantic relations between the words at either end of a link. There are four main categories of methods for determining the characteristics of semantic relations that have been applied to developmental questions: (1) methods based on expert knowledge, in the form of a dictionary or thesaurus (Steyvers & Tenenbaum, 2005), (2) methods based on word association norms (Hills et al., 2009a, 2009b; Steyvers & Tenenbaum, 2005), (3) methods based on the statistical analysis of a corpus (Beckage, Smith, & Hills, 2011; Hills 2013; Hills et al., 2010) and (4) definitional methods based on empirical data, in the form of semantic feature norms (Beckage, Aguilar, & Colunga, 2015; Engelthaler & Hills, 2016; Hills et al., 2009a, 2009b; Yurovsky, Bion, Smith, & Fernald, 2012). While these methods all result in some index of semantic relatedness, the patterns and magnitudes of relations and the variance they account for in empirical data vary (De Deyne, Peirsman, & Storms, 2009; Maki & Buchanan, 2008; Mirman & Magnuson, 2007; Riordan & Jones, 2011). This variance indicates the different methods either capture different aspects of semantic relations (Riordan & Jones, 2011), or differences in the activities resulting in the products used by the methods constrain semantic relations in different ways (Maki & Buchanan, 2008). Either way, the choice of link type plays an important role in the types of research questions that can be approached.

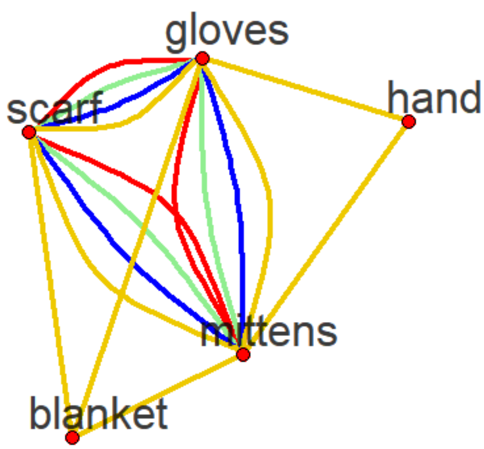

Subset of Features for GLOVES

Feature	Type
comes in pairs	Encyclopedic
worn in winter	Encyclopedic
worn for warmth	Functional
worn for surgery	Functional
made of wool	Perceptual
has fingers	Perceptual
has a thumb	Perceptual
clothing	Taxonomic

Figure 1. A noun-feature network containing four neighbors of GLOVES (left) and a subset of the features of GLOVES.

While each of the methods has its strengths, this dissertation focuses on graph-theoretic networks defined using semantic feature norms because of (1) the clear alignment with the proposed mechanism of feature-guided selective attention and, relatedly, (2) how they enable the simultaneous characterization of both which concepts are connected and how they are connected. The second point is demonstrated in figure 1 by the different types of features making up the substructure of the links between GLOVES and its neighbors. To define the links researchers have mostly used a set of feature production norms collected by McRae and colleagues (2005). The features making up this set have also been categorized based on a brain-based knowledge type taxonomy that divides semantic features into four main types (Cree & McRae, 2003). *Perceptual features* describe information that is accessible through one of the five senses (e.g., has a round end, made of metal, etc.). *Functional features* describe the ways in which people interact with objects (e.g., used for eating, used for stirring, etc.). *Taxonomic features* refer to positions within a system of classification (e.g., a utensil). Finally, there are a number of facts that don't fall neatly into any of these categories (e.g., found in kitchens, part of a table setting, etc.), these are all grouped together as *encyclopedic features*. These different feature types have played a central role in the questions leading researchers to use this methodology, and are an important part of why we used this methodology to explore feature biases as a potential explanation for the evidence of relations between toddlers' word meaning knowledge and word processing abilities.

Characterizing noun-feature network structure

The primary reason researchers use graph-theoretic methods for questions involving word knowledge is that there are numerous ways of characterizing patterns of semantic relations between words for multiple levels of analysis, including neighborhood-, and lexicon-level structure. Furthermore, for each level of analysis, metrics can be calculated that are either centered on individual noun-concepts or describe the network as a whole.

Neighborhood-level metrics describe the structural characteristics of the immediate neighborhood of a word. Two commonly used metrics are: Degree and Clustering Coefficient. *Degree* is defined as the number of neighbors that a word has, and has been emphasized in accounts of semantic richness (Pexman et al., 2008), word processing (Mirman & Magnuson, 2008), and word learning (Storkel, 2009; Storkel & Adlof, 2009). *Clustering coefficient* quantifies the degree to which a word's neighbors are linked, and is calculated by taking the ratio of the number of links between neighbors over the maximum possible number of links between neighbors. Clustering coefficient ranges from 0, when none of a word's neighbors are connected, to 1, when all of a word's neighbors are directly connected. There is building empirical and modeling evidence that clustering coefficient relates to patterns of inhibition and facilitation in noun-concept processing contingent on activation feedback between neighbors (Mirman & Magnuson, 2008; Vitevitch, Ercal, & Adagarla, 2011).

Lexicon-level metrics characterize of word in the context of the entire lexicon. The most basic lexicon-level metric is *distance* (aka *path length*), defined as the minimum number of links that must be traversed to get from one node to another. The mean distance for a given node relates to how centrally located it is within a network, and may have functional value related to the efficient spread of activation across the full range of the network, aiding fluent processing and word retrieval (Griffiths, Steyvers, & Firl, 2007; Steyvers & Tenenbaum, 2005).

In addition to higher-level structure, noun-feature networks provide the unique ability to simultaneously characterize local-level structure that *only* considers immediately local information of the noun-concept. One such metric, *Number of Features* (NoF), is calculated as the number of semantic features a concept has. NoF has been argued to characterize the semantic richness of a concept (Pexman, Lupker, & Hino, 2002; Pexman, et al., 2008), and relates to the speed and accuracy of adult performance in lexical decision and semantic categorization tasks (Grondin, Lupker, & McRae, 2009; Pexman et al., 2002; 2008), and potentially relates to the ease of word learning (Rabovsky et al., 2012).

The interpretations for each of the metrics described above are based on the idea that graph-theoretic noun-feature networks to some extent capture the structure of semantic relations that drive aspects of lexico-semantic processing noted in section 2.1. However, a narrower interpretation that can be applied to noun-feature networks created using children's word knowledge is that structure emerges out of the patterns of pairwise relations that influence the process of learning each word. In other words, the structure is a signal of the dynamic system underlying the word learning processes that led to each word entering the network, rather than semantic cognition more broadly. This distinction is subtle, and the two interpretations are non-exclusive, but while the former – semantic structure interpretation – comes with extensive theoretical baggage regarding the mechanisms underlying semantic cognition, the latter – signal of learning dynamics interpretation – remains more agnostic regarding such mechanisms. One constraint of taking the latter approach is that the interpretations for different metrics are more closely tied to the specific research question at hand. Regardless, much of the research described in the following section is framed within the signal of learning dynamics interpretation of noun-feature based lexico-semantic networks, and this is also the interpretation taken in the empirical studies described in chapters 2 and 3 of this dissertation.

The Semantic Feature Structure of Early Word Knowledge

In this section we review research that applies the graph-theoretic noun-feature network methodology to explore early language development, and discuss how the work informs the central hypothesis of this dissertation (i.e., noun-concept processing skill relates to the structure of a toddler's vocabulary via selective attentional biases driven by features shared between to-be-processed words and known vocabulary). We conclude with a recent project demonstrating perceptual features matter most in the mind of a toddler – setting the stage for the experiments described in chapters 2 and 3.

Using graph-theoretic noun-feature networks to explore lexico-semantic development

Researchers using noun-feature networks to explore developmental questions have typically used nodes consisting of the concrete nouns on the MacArthur-Bates Communicative Developmental Inventory, Words and Sentences form (MBCDI:WS; Fenson et al., 1996), a parental checklist of words commonly produced by 16 to 30-month-olds. Though the measure is an incomplete sample of potential word-knowledge, it has repeatedly related with a variety of group and individual differences in lexical and lexico-semantic processing (e.g., Borovsky et al., 2016; Fernald, Perfors, & Marchman, 2006; Rämä, Sirri, & Serres, 2013). Also, as noted in the previous section, the capability to categorize features into different types is often the key reason researchers use this methodology, with the focus generally being on perceptual versus functional features, while excluding taxonomic and encyclopedic features – which are argued to be inaccessible by young children (e.g., Hills et al., 2009a).

Hills and colleagues (2009a) compared noun-feature networks formed via either perceptual or functional features to explore how features of each type drive early lexico-semantic network structure. They found that both kinds of noun-feature networks contained clusters of words that resemble adult superordinate taxonomic designations (e.g., animals, clothes, food, toys, etc.). However, perceptual networks contained more links and had a higher percentage of links made up of multiple correlated features – for example, how things that *have wings* are also likely to *have beaks* and *have feathers*. Such feature correlations have been shown to influence adult performance in semantic tasks (McRae et al., 1997; McRae, Cree, Westmacott, & De Sa, 1999), so structural patterns in early noun-feature networks that capture such correlations might relate to the ease with which different words are learned. Indeed, these results directly relate to the proposal that shared-feature-driven selective attention is underlying the facilitatory effects of high category density, and furthermore highlight perceptual features as being of particular interest.

Yurovsky and colleagues (Yurovsky, Bion, Smith, & Fernald, 2012) compared how the structure of toddlers' perceptual versus functional noun-feature networks related to their use of Mutual Exclusivity – the process of assigning a novel label to a novel object, when presented with a novel object in the context of known objects (Markman & Wachtel, 1988). They hypothesized that only functional structure should relate to the use of mutual exclusivity, because over the course of the toddlers' lives mutual exclusivity should most reliably have aided learning for objects that co-occurred together in the environment – and thus are likely to share functional features but not necessarily perceptual features. This proved to be the case, with mutual exclusivity use being correlated with functional, but not perceptual, noun-feature network structure. Interestingly, by discussing Mutual Exclusivity in the context of shared functional features, this work places the mechanism in the same framework as feature biases. Furthermore, by using functional features as a way to capture the co-occurrence of objects in situations, this work highlights the overlap between the noun-feature approach and methods based on the co-occurrence of noun-concepts in a corpus. Indeed, comparing patterns of concept clustering using both feature- and corpus co-occurrence-based models of semantic representations, Riordan and Jones (2011) found that the information provided by functional features is redundantly coded in linguistic experience, while perceptual features provide unique information. The reason this point matters for this dissertation is that the neural network models that duplicated the facilitatory effect of high-density categories on word processing had semantic representations that are entirely based on patterns of co-occurrence in the corpus used to train the networks (Borovsky & Elman, 2006). Thus, while the Attentional Learning Account and above work by

Hills and colleagues highlight the role of perceptual features, this work by Yurovsky and colleagues (2012) and the links with co-occurrence techniques indicate functional features may also be worth looking into. While we don't follow this path in the current project, in the concluding chapter we note this is a promising avenue for future work

Finally, Engelthaler and Hills (2017) also explored whether noun-feature networks capture evidence that early word learning is driven by mutual exclusivity. However, they proposed the mechanism extends beyond the traditional definition to include situations where a child is trying to determine whether a novel object belongs to a named category. Based on this extension, they hypothesized that early in development more distinctive noun-concepts would be more likely to be differentiated and receive object-labels, and thus would be learned earlier. Accordingly, rather than using network metrics based purely on shared features, as is typically done with noun-feature networks, they used metrics that also consider differences in the featural make-up of words. Using such measures of noun-concept distinctiveness, they found that early learned words are indeed more distinct, particularly in regard to their perceptual featural makeup. This work is worth considering precisely because it provides a counterpoint to the dominant approach based solely on shared noun-features. Furthermore, it completes the linkage of mutual exclusivity and feature biases, collapsing them into a single mechanism and requiring an update to the earlier description of feature-driven shifts in attentional selection to include a potentially key role for feature novelty.

The above research demonstrates how noun-feature networks can be used to advance our broad understanding of lexico-semantic development and provide information that is specifically relevant to the idea that shared-feature driven shifts in selective attention may influence early word learning and processing. However, two key limitations constrain the generalizability of these results. First, prior analyses used feature norm data sets (e.g., McRae et al., 2005) that miss many early learned words. This limited Hills and colleagues (2009b) to exploring networks using only 130 of the 359 early learned nouns on the MBCDI:WS and led Engelthaler and Hills (2017) to use words that young children are unlikely to know in conjunction with adult generated Age of Acquisition values. Second, while the reviewed research explored metrics that characterized network structure over both small and large grain sizes, they did not consider whether their findings could be driven by even more granular measures that only consider information that is immediately local to either nodes or links. Importantly, the level of structure at which any effects originate has important implications for our broader understanding of early lexico-semantic processing and may also offer insights regarding the mechanism that is of specific interest to this dissertation: shared-feature-driven shifts in selective attention. In regard to broader implications, modeling evidence suggests that large-scale structural patterns facilitate the spread of activation across the lexicon, supporting fluent word processing and retrieval (Griffiths, Steyvers, & Firl, 2007; Steyvers & Tenenbaum, 2005). However, empirical evidence suggests that more local-level structure also facilitates word processing and learning (Borovsky & Peters, submitted), and even more granular measures that characterize the features of a word could elucidate how the semantic properties of a word itself drives early learning, and relates to higher-level patterns. In regard to shared-feature-driven shifts in selective attention, the structural locus of any effects may serve to highlight the relative roles of either the shared-feature(s) or the features to which attention is shifted to. In this next section, we describe a recent project that overcomes these two limitations and serves as the groundwork for the empirical studies presented in chapters 2 and 3.

Peters & Borovsky (2019): Perceptual Features Matter Most

Although it's widely agreed that adults are able to flexibly encode the full variety of meanings outlined by the four feature types (i.e., encyclopedic, functional, perceptual and taxonomic), there is an ongoing debate regarding the developmental timecourse of when different types of meaning become accessible early in life. On one side of the debate researchers highlight the early importance of functional features. For example, the *thematic-to-taxonomic shift* (e.g., Inhelder & Piaget, 1964; Smiley & Brown, 1979) proposes that thematic (aka functional or associative) relations underly early semantic processing, while other types of meaning used in determining taxonomic relatedness are learned and prioritized later. On the other side of the debate are those who argue for a *perceptual-to-conceptual* shift (Quinn & Eimas, 1997), in which early categorization and word learning depend on perceptual features of objects – while more abstract, conceptual information is only represented after the lexico-semantic system develops sufficiently. Peters and Borovksy (2019) sought to shed light on this debate with two modeling studies that explored the patterns of early word learning using noun-feature networks in combination with a newly developed extension of the McRae feature norms to allow for consideration of all 359 nouns on the MBCDI:WS (Peters, McRae, & Borovsky, in prep).

Study 1 sought to explore how the featural characteristics of word meanings contributed to the order of early word learning. The approach involved using multiple regression to model words' empirically derived Age of Acquisition (AoA) – calculated using over 5000 administrations of the American English version of the MBCDI, accessed from Wordbank (Frank, Braginsky, Yurovsky, & Marchman, 2017) – as a function of their semantic feature characteristics, while controlling for word frequency in Child Directed Speech. Feature characteristics were varied along two dimensions: (1) feature type: encyclopedic, functional, perceptual, and taxonomic; and (2) level of structure: local (word)-, neighborhood-, and lexicon-levels. Results indicated that perceptual feature characteristics were the most robust predictors of AoA across all levels of structure – with words having more perceptual features, more perceptual links, and perceptual links bringing them closer to the rest of the network, all being learned earlier according to analyses for each of the three levels of the structure. Furthermore, while model fits were significant across all levels, the local-level model was notably better than the neighborhood- and lexicon-level models. One limitation with the analytical approach of study 1 was the inability to explore potential non-linear relations between AoA and feature characteristics. This problem was exacerbated by the fact that, following precedent in the literature, network metrics for words were calculated using the full lexicon, rather than the subset of words that would have made up the lexicon at the point in development when the word was learned. Thus, study 2 used a novel method that avoided this limitation.

Contrasting with study 1, in study 2, lexico-semantic development was explicitly modeled. Words were entered into a network following their AoA rank order, and the semantic feature characteristics of the network were calculated for each time step, resulting in normative growth trajectories for each of the metrics. These normative trajectories were then compared to random trajectories for growth networks in which words were entered into the networks in a random order. This comparison of normative versus random trajectories allowed for the exploration of longitudinal differences that resulted precisely because of the order in which words are normatively learned. Following study 1, analyses varied along the two dimensions of feature type and level of structure. Results paralleled those of study 1, indicating that perceptual featural characteristics related to the normative order of early word learning for all levels of structure.

Broadly, the findings of Peters and Borovksy (2019) suggest that perceptual features and links between noun-concepts support early word learning. Local-level findings across the two studies indicated that words with more perceptual features are, on average, learned earlier. This pattern is consistent with evidence that concepts with greater concreteness and iconicity are learned earlier (e.g., Gilhooly & Logie, 1980; Perry et al., 2015) and work indicating that objects with greater perceptual salience draw attention, facilitating subsequent learning of object-label pairings (Pruden et al., 2006). Neighborhood-level findings across the two studies also indicated that words with more perceptual linkages to other words are learned earlier. This pattern of results is very similar to findings that adult conceptual processing is facilitated for concepts with more perceptual features, and in particular more perceptual features that are shared with other concepts (Grondin et al., 2009). However, in study 1, the neighborhood-level effect was weaker than at the local-level, and in study 2, the neighborhood-level period of systematic difference between the normative and random growth trajectories was nearly identical to the local-level. Together, these results raise the question of whether the neighborhood-level effects are actually the result of a robust lower-level effect propagating up to higher levels. In other words, higher-level structure might not actually play a driving role in determining the order of early word learning and lexico-semantic network structure, but instead arise out of the ease of individual word learning. If this is true, this arguably calls into question the entire mechanism of shared-feature-guided shifts in selective attention as a possible underlying driver of early word learning, or at the very least shifts the importance in the mechanism towards the features to which attention is shifted. However, it is important to note that Peters and Borovsky (2019) modeled *normative* patterns in early lexical development, while the Attentional Learning Account (Smith & Samuelson, 2006) was specifically designed to explain patterns in early word learning biases in an individual's experiences. In other words, although this work usefully highlights the relative importance of perceptual features, it likely obscures patterns in development that emerge from and result in individual differences.

Conclusion

In this introductory chapter, we first reviewed research exploring how word meaning knowledge influences early word processing and learning capabilities. To explain the results, we introduced the idea of shared-feature-guided selective attention (Jones et al., 1991; Jones & Smith, 1998) as an application of the Attentional Learning Account (Smith & Samuelson, 2006) to describe idiosyncratic attentional biases that could both emerge from and result in clusters of known words united by shared features in individual toddlers' vocabularies. We then introduced feature-based lexico-semantic networks as a useful tool for characterizing the structure of toddlers' productive vocabularies, and reviewed research that applies the methodology to explore early language development. Finally, we concluded with a recent project demonstrating perceptual features matter most in the mind of a toddler. Altogether, the work presented here demonstrates the utility of noun-feature networks for exploring developmental questions, and clearly lays the groundwork for a way forward. Namely, empirical research is necessary to explore whether shared perceptual features truly influence early word processing, and if so, whether shared-feature-driven shifts in selective attention are a possible mechanism. The remainder of this dissertation presents two eye-tracking studies that empirically explore whether early word processing and learning is indeed facilitated by shared perceptual features with other concepts. In chapter 2, we present work investigating how patterns of perceptual connectivity in the productive vocabularies of toddlers influence their initial attentional biases to known objects and subsequent label processing. In chapter 3, we present the results of a word-learning task likewise designed to investigate how patterns of shared perceptual features influence attentional biases to novel objects and subsequent learning outcomes of object-label pairings. Finally, in chapter 4, we conclude by discussing the findings of the two empirical studies in the broader context of individual differences in early word learning and point the way forward.

CHAPTER 2. PERCEPTUAL-SEMANTIC CONNECTIVITY INFLUENCES ATTENTION TO KNOWN OBJECTS AND SUBSEQUENT LABEL PROCESSING

Introduction

Recent modeling work has demonstrated that, at the *group* level, patterns of shared perceptual features between early-learned noun-concepts drive lexico-semantic development in 16- to 30-month-olds (Peters & Borovsky, 2019). A natural follow-up question asks: Do patterns of perceptual similarity within *individual* toddlers' noun-concept vocabularies likewise influence online lexico-semantic processing? And if so, what is the underlying mechanism? Recent work separately demonstrating the semantic structure within individual children's early productive vocabularies relates to word processing skill on the one hand (e.g., Borovsky, Ellis, Evans, & Elman, 2016) and their patterns of attention to perceptual features of objects on the other (e.g., Colunga & Sims, 2017), hints at an interesting possibility. Namely, noun-concept processing may relate to the semantic connectivity of a toddler's vocabulary via attentional biases that depend on relations between the semantic features of to-be-processed words and known vocabulary. Here, we build on this prior work and ask whether toddlers attention to known objects and their labels is influenced by objects' patterns of perceptual connectivity with other noun-concepts in individual toddlers' productive vocabularies.

Conceptual development involves the recognition of both individual concept features and patterns of relations between concepts. The integration of such semantic structure into young children's lexico-semantic systems allows them to form categories (e.g., Quinn & Eimas, 1996), which in turn can be leveraged to learn new words. Borovsky and colleagues (2015) found that 2-year-old toddlers are better able to learn words for new category members when they know a higher proportion of items in the category. Furthermore, they found this facilitatory effect of high category density extended to known word processing (Borovsky, Ellis, Evans, & Elman, 2016). Neural network simulations (Borovsky & Elman, 2006) provided a potential explanation: lexico-semantic systems with more interconnected semantic structure are more easily able to recognize and map similarities from known to novel items, and in turn are more easily able to integrate novel items. Notably, the initial step in the process – recognizing shared features – approximately maps onto a well-studied mechanism in early word learning: the selective attention to relevant object features when generalizing names to new objects (Smith, Jones, & Landau, 1988).

According to the Attentional Learning Account (e.g., Smith & Colunga, 2008; Smith & Samuelson, 2006), children learn to selectively pay attention to features relevant to different domains of objects through experience. For example, for many count nouns, such as SPOON, object shape in and of itself is a defining feature. However, some count nouns, such as those with eyes or feet, may necessitate attending to a wider set of features for accurate recognition and generalization (Jones, Smith, & Landau, 1991; Jones & Smith, 1998). Over the course of learning a variety of noun-concepts, children attune to the coherent structure across linguistic and nonlinguistic input in naming situations, and then (via associative learning) heuristically capitalize on this structure when learning words and selectively attend to object features that have correlated with similar contexts in their experience (Smith et al., 2002). Importantly, this process depends on the integration of such structure into the lexico-semantic system of the individual, demonstrated through differences in usage of selective attentional biases related to individual differences in vocabulary composition (Colunga & Sims, 2017; Jones, 2003; Perry & Samuelson, 2011). This is particularly true for children at the lower end of the vocabulary size spectrum, who have vocabularies with more variable composition than children with larger vocabularies.

The Attentional Learning Account can be used to expand on the explanation for the facilitatory effect of high category density on early word learning and processing. Children who know a variety of items in a given category have more experience with the structure of coherent covariation (Rogers & McClelland, 2004) of features across items in the category than children who know fewer items. In turn, when encountering novel items, high category density may enable children to not only more easily recognize shared features, but also employ feature-guided selective attention to relevant feature domains based on their knowledge of the correlational structure of the category. Such feature-guided selective attention has been argued to jump start the creation of cohesive representations (Jones et al., 1991), which in turn could be beneficial for subsequent object-label mapping. Notably, the benefits of this process could extend beyond the initial name learning episode, snowballing across subsequent naming interactions with the object and resulting in even richer, more cohesive, and more easily accessible representations. Furthermore, considering word learning is an ongoing process that continues well past the first exposure and relates to improvements in familiar word recognition (McMurray, Horst, & Samuelson, 2012), this explanation cleanly ties together the facilitatory effect of high category density for both novel and known object recognition. Specifically, novel objects in high-density categories may benefit from an initial feature-guided selective attention derived processing boost, while known objects in high-density categories may have a history of such facilitated processing – which in turn has pushed the words further along the familiarity continuum then known words in low-density categories, resulting in relatively facilitated processing.

Rather than a broadly applicable mechanism, such shared-feature-guided processes should be more idiosyncratic to the clusters of known objects that form categories. To explore this issue in greater depth, we need to move beyond simple descriptions of categories or feature domains (i.e., shape), and work with more granular featural descriptions of objects to explore patterns of shared features across the full breadth of individual toddlers' vocabularies. Therefore, in the following section we introduce the use of noun-feature networks to model toddlers' vocabularies – an approach that builds on recent advances in the application of graph-theoretic network modeling techniques to lexico-semantic representations (e.g., Hills et al., 2009a, 2009b; Steyvers, & Tenenbaum, 2005).

Using noun-feature networks to characterize vocabulary structure

Noun-feature networks can be visualized as a set of nodes and links. The nodes represent noun-concepts, and the links represent shared features according to a set of feature production norms (e.g., McRae et al., 2005). Feature production norms are collected by presenting participants with a list of concepts, and asking them to write down the features that come to mind. Concepts are given to multiple participants, and only features that are produced repeatedly make it into the final set of norms.

Features can be categorized into several possible types (e.g., Cree & McRae, 2003), of which functional and perceptual features have been of particular interest to developmental researchers. *Functional features* refer to the ways people interact with objects (e.g., used for eating), while *perceptual features* refer to information that is available through one of the five senses (e.g., has a round end). These feature types can be used to decompose networks to separately explore how different types of features drive early lexico-semantic development. Using this technique, Hills and colleagues (2009a) found that perceptual and functional features both result in networks containing clusters that resemble adult superordinate taxonomic designations (e.g., animals, clothes, food, etc.), but perceptual networks contained more links and had a higher percentage of links made up of multiple correlated features. Importantly, this result directly relates to the above proposal that shared feature driven selective attention is underlying the facilitatory effects of high category density, and highlights perceptual features as being of particular interest.

The relative importance of perceptual features is further supported by work exploring how the order in which early words are learned relates to their featural characteristics (Peters & Borovsky, 2019). Broadly, the findings suggest that perceptual features and links between noun-concepts are the most important type of feature supporting early word learning. Furthermore, in supplementary analyses they discovered the results were driven by two perceptual feature sub-types: visual-motion and visual-form and surface features. Building on this work, here we focus on these two perceptual feature sub-types, and ask whether toddlers' attention to known objects and their labels is influenced by the objects' patterns of (visual-motion and visual-form and surface) perceptual connectivity with other noun-concepts in individual toddlers' vocabularies.

Using eye-tracking to index online object and label processing

To explore the online processing of objects and their labels we use a variant of the *Looking While Listening Paradigm* (LWLP; e.g., Fernald et al., 2008) which indexes online processing via measurements of eye gaze to visual images of objects accompanied by their spoken labels. This variant was used in the aforementioned work exploring the effects of category density on novel and known word processing (Borovsky et al., 2015; 2016). Importantly, this paradigm is specifically designed to explore gaze patterns made in response to spoken stimuli, and experiments are constructed so as to counterbalance out any potential effects of pre-labeling object attentional biases. While proper counterbalancing does mean that initial biases that align with the target in one case will align with the distractor in another, it does not actually allow for differentiating effects that wholly result in the course of label processing from those that involve cascading effects from pre-labeling biases. The predecessor of the LWLP, the *Intermodal Preferential Looking Paradigm* (Golinkoff, Hirsh-Pasek, Cauley, & Gordon, 1987), takes a different approach by including a 'salience trial' before any test or learning trials, which can then be used to either check that mean fixation proportions for a pair of stimuli are equivalent or correct mean fixation proportions in test trials when biases do exist (e.g., Naigles, 1990). Importantly, the initial salience of an object for an individual is possibly related to the perceptual knowledge of the object they bring to the task. Thus, we take a combined approach making use of both counterbalancing and explorations of attentional biases in the initial preview period and their influence on fixations during the test period.

Participant differences that may relate to variation in feature-guided selective attention

Feature-guided selective attention is proposed to emerge via a developmental process that unfolds over the course of a variety of early experiences with objects and language. As such, individual and age-related differences that influence or constrain such experiences should relate to variation in patterns of feature-guided selective attention. In this project, in addition to age, we consider two potentially relevant individual differences: word learning skill and temperamental tendency to maintain focused attention.

The age of an individual constrains their history of language learning experiences in a variety of ways, including the range and variety of potential experiences they could have been exposed to – with younger individuals on average having been exposed to fewer and less varied experiences – and the degree to which the contents of any given experience were able to make contact with the individual's learning system – with younger individuals, with less developed memory and attentional systems and less word knowledge, having less access to the full depth of potentially learnable content in any given learning opportunity (e.g., Rakison, Lupyan, Oakes, & Walker-Andrews, 2008). Such constraints on learning histories would influence the patterns of correlational structure in lexico-semantic knowledge that are argued to drive shared-feature guided selective attention. Furthermore, evidence that the influence of semantic relations on processing develops over the second life (e.g., Rämä, Sirri, Serres, 2013) suggests age may also relate to differences in the online usage of shared feature guided selective attention.

Individual differences in word learning skill and tendency to maintain focused attention might similarly constrain and shape the patterns of correlational structures in word knowledge over time and/or influence such a mechanism online. Children who are better word learners, and have more knowledge than their lower skill peers, will have wider access to the depth of content in a learning opportunity. While word knowledge also varies with age, if word learning skill is measured in comparison to same-age peers – such as via vocabulary size percentiles – we can explore variance that is orthogonal to age. Finally, children who tend to maintain focused attention in learning opportunities, should have a broader sampling of co-occurrence data compared to their peers who pay less attention in such situations (for a related discussion see Smith, Suanda, & Yu, 2014). While a variety of attentional mechanisms likely play a role in early word learning, here we focus on attentional focusing given evidence that it relates to individual differences in word learning outcomes (Dixon & Shore, 1997; Dixon & Smith, 2000). Thus, age, word learning skill, and tendency to maintain focused attention, may all relate to individual differences in shared-feature-guided selective attention in similar, but dissociable ways.

The current study

The primary goal of this project was to explore how the (visual-motion and visual-form and surface) perceptual-semantic connectivity of known objects influences both patterns of pre-labeling attentional biases and subsequent object-label processing. A secondary goal was to determine whether any semantic connectivity effects are related to differences in age, word learning skill, or tendencies to maintain focused attention. To test these questions, we exposed 24- to 30-month old toddlers to known object-labels pairs that vary in patterns of (visual-motion and visual-form and surface) perceptual-semantic connectivity based on a set of feature production norms for early learned nouns (Peters, McRae, & Borovsky, in prep). As described in further detail in methods, we use measures of connectivity based on degree: the number of links a noun-concept has. As links are defined based on shared features, words with higher degree have more shared features with other words, and thus more potential cues for shared-feature-guided selective attention to relevant aspects of correlational word knowledge. We measured pre- and post-labeling attention to object pairs using the aforementioned variant of the LWLP (Borovksy et al., 2015; 2016), and characterized attentional biases and label-processing based on the proportion of looking time to each item in a pair.

Thus, the central hypotheses driving this work are that the semantic connectivity of known objects will influence (1) toddlers' patters of attention and (2) subsequent patterns of object-label processing. We consider four possible outcomes, each of which have different implications for mechanism underlying how attention and connectivity support learning:

1. Attentional biases drive looks to objects with *high* connectivity, and support subsequent processing of labels for those objects.

2. Attentional biases drive looks to objects with *low* connectivity, and support subsequent processing for those objects.
3. *High* semantic connectivity supports processing of object labels, irrespective of attentional biases during preceding free-viewing opportunities.
4. *Low* semantic connectivity supports processing of object labels, irrespective of attentional biases during preceding free-viewing opportunities.

Outcome 1 is the predicted outcome, given it follows naturally from the above explanation for the facilitative effects of high category density on label processing based on the Attentional Learning Account, assuming the effect derived from pre-labeling effects and will remain even in the presence of potential semantic interference resulting from being paired with a same category item. Conversely, outcome 2 could result from some kind of distinctiveness or novelty effect, or from inhibited processes of the high connectivity item due to semantic interference. Both of these first two outcomes – in which patterns in the test period parallel those of the preview period – would lend support to the hypothesis that connectivity driven attentional biases towards objects in the environment influence known object-label processing and, following the chain of logic that known and novel words lie on a continuum, played a role in the preceding label learning experiences that resulted in the current state of the object-label representation. In contrast, outcomes 3 and 4 – in which processing patterns in the test period are unrelated to attentional biases in the preview period – would indicate that, in the context of same category items, semantic connectivity influences attentional biases and label processing in different ways that outweigh any potential cascading effects from the former to the latter.

Methods

Participants

Sixty 24- to 30-month-old toddlers were recruited via outreach events and community fliers. We chose this age range because it is a time of particularly dramatic change in children's lexico-semantic systems. However, we did not include lower ages because as productive vocabulary sizes get smaller so too do differences in patterns of connectivity – making it difficult to choose appropriately balanced stimuli. Six children completed the study but were not included in analyses for the following reasons: history of chronic ear infections and hearing loss concerns (n=2), neurological impairment (n=1), insufficient data due to track-loss/fussiness (n=3). This resulted in 54 toddlers being included in the final analyses (26 F, 28 M), all of who were reported to meet the following pre-registered (see Appendix C) inclusionary criteria: normal vision, hearing, and no history of chronic ear-infections; no history of neurological or cognitive impairment; not born preterm (<37 weeks) and low birth weight (<5lb8oz); not exposed to language other than English for more than 8 hours a week (~10% of waking hours).

Approach to calculating metrics of semantic connectivity

Features

This work uses an extension (Peters, McRae, & Borovsky, in prep) of the McRae semantic feature production norms (McRae et al., 2005), which now include all 359 noun concepts on the MBCDI:WS. The features are classified into four broad types (encyclopedic, functional, perceptual, and taxonomic) according to the Cree and McRae (2003) Brain region knowledge type taxonomy. Perceptual features can be further broken up into an additional seven subtypes: smell, sound, tactile, taste, visual-color, visual-form and surface, and visual-motion.

Only two types of perceptual features – visual-form and surface and visual-motion features – are used in this work based on recent modeling evidence they are the primary types of features related to the ease of early word learning (Peters & Borovsky, 2019).

Network construction

We constructed graph-theoretic lexico-semantic networks for each individual, based on their productive noun vocabulary according to the MBCDI:WS, in which nodes represent known nouns and links between nodes represent shared perceptual (specifically: visual-form and surface and visual-motion) features. Links were treated as unweighted, undirected edges, following prior precedent (e.g., Engelthaler & Hills, 2017; Hills et al., 2009a).

Using the same method, we also created a normative 26-month-old's network containing all nouns with an Age of Acquisition (AoA) less than or equal to 26 months. We calculated AoA for each of the 359 concrete nouns on the MBCDI:WS using the Wordbank database (Frank et al., 2017) – an open repository of MBCDI data, of which we used 5450 administrations of the MBCDI:WS. Following Braginksy, Yurovsky, Marchman and Frank (2016), for each word we calculated the proportion of children that produced the word for each of the ages (ranging from 16 to 30 months. We then fit the logistic curves to the proportion data, and took the age at which the curves crossed .5 as our age of acquisition.

Calculating metrics of semantic connectivity

This project focuses on a set of measures based around one of the simplest network metrics: degree. The *degree* of a word is simply defined as the number of links it has to other words in the network. Given links are defined using shared features, degree characterizes the pattern of shared features between a word and its neighbors and is therefore closely related to the research questions driving this work. The three metrics we use are: normative degree category, corrected degree, and corrected degree difference. We calculated the *normative degree* of each item as its degree in the normative 26-month-old's (visual-form and surface and visual-motion) perceptual network described above. As described in detail in the next section, paired items were rank ordered in terms of normative degree, and assigned to either "high" or "low" *normative degree categories*. For each participant, we also used the networks created using their individual productive noun vocabularies, to calculate the *corrected degree* (a.k.a. normalized degree) for each item, defined as the degree of the item divided by maximum possible degree for that individual's network (i.e., the number of words in the network minus one). We used corrected degree instead of degree because the degree of items in an individual's network are closely related to the number of items in that network, which in turn is closely related to the age and MBCDI:WS percentile of the individual. In appendix A we present analyses validating this problem with degree is alleviated by using corrected degree. Finally, for each pair of items we calculated the *corrected degree difference* between normative high and low items as the corrected degree of the high item minus the corrected degree of the low item.

Materials

Item selection

Words were chosen from the animal and vehicle categories in the MBCDI: Words and Sentences. These words were presented as part of a larger study that included novel objects. The novel objects were assigned perceptual features that typically go with either animals or vehicles, so the categories for the known words were chosen so as to match with the novel objects.

We used the Wordbank database to select six highly known words in each category (i.e. produced by at least 75% of 26-month-olds) for a total of 12 words. The items were then organized into yoked pairs, with three pairs of animals (bug-chicken, bear-duck, cat-fish) and three pairs of vehicles (bus-boat, truck-bike, airplane-train). Pairs were selected so that within a given pair one item had a high perceptual degree, calculated using the normative 26-month-old's network, and one had a low perceptual degree (see Figure 2).

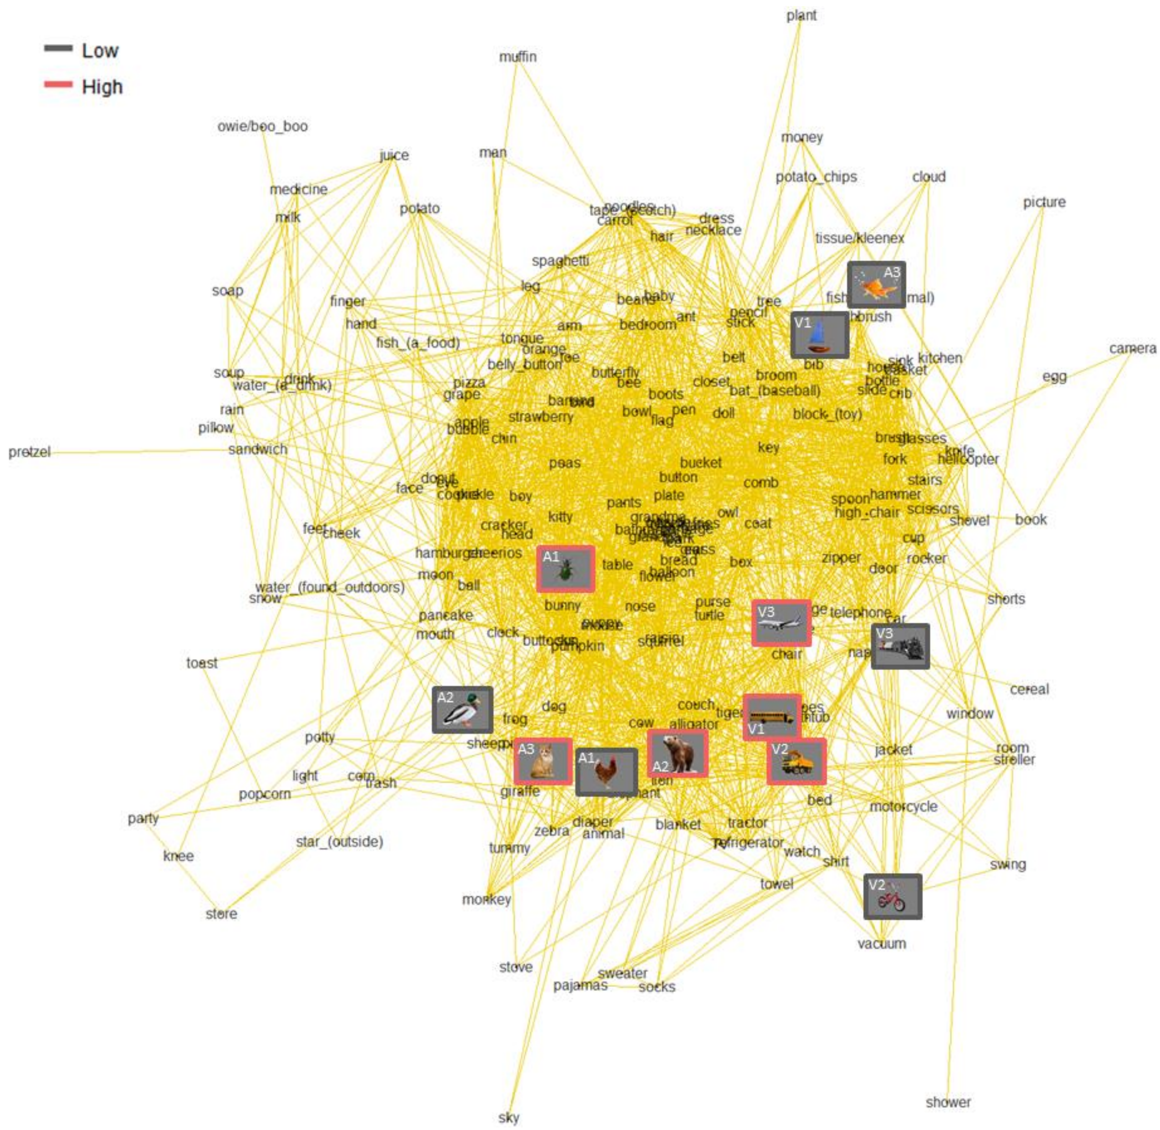

Figure 2. Main component of a normative 26-month-old's (visual-motion and visual-form and surface) perceptual network, containing 219 noun-concepts. Images for selected low (grey) and high (red) normative connectivity stimuli are shown. Pairs are marked with white text, with numbering matching the order of items presented in Table 16A, in Appendix B.

Furthermore, the yoked pairs were chosen to control for a range of other factors that have potential influence on patterns of psycholinguistic processing. The factors that we controlled for were potentially relevant variables, including: total degree (i.e. degree in a network with connections determined using all features), number of perceptual features, number of visual-form and surface and visual-motion features (with toddler accessibility above the median), CHILDES word frequency, number of phonemes, number of syllables, bigram frequency, phonological neighborhood density, AoA and proportion of children producing the word at 26-months of age. T-tests also confirmed that all controlled factors were not significantly different across pairs (see Appendix B for details).

Auditory stimuli

A female Standard American English speaker spoke in an infant-directed voice for all auditory experimental stimuli and additional encouraging phrases. The spoken stimuli were recorded on a mono channel at 44.1 kHz sampling rate. All experimental stimuli were adjusted to a mean duration of 1016 ms, and all stimuli, including encouraging phrases and an attention getter, were adjusted to a mean intensity of 70 dB. The encouraging phrases were used to help keep the toddlers' interested in the experiment, and included phrases such as "Isn't that cool?", "Do you like it?", and "Isn't this fun!" The attention getter consisted of a recording of the speaker saying "Look" at the onset of the gaze-contingent center stimulus that appeared in each experimental trial before the onset of the spoken experimental stimulus.

Visual stimuli

Visual experimental stimuli consisted of 600x450-pixel photorealistic color images on a 1920x1080-pixel screen. Object images were selected so as to match toddlers' experiences with the experimental items, verified by consulting with parents and laboratory members. All images were isolated on a grey background. Figure 3 shows an example of the visual stimuli in an experimental trial. Other images were also used to keep and direct the attention of children throughout the study. These images included large pictures of characters from popular children television shows (e.g., Elmo, Paw Patrol characters, etc.) and small images used to direct children's attention to the center of the screen (e.g., smiling stars, happy faces, etc.).

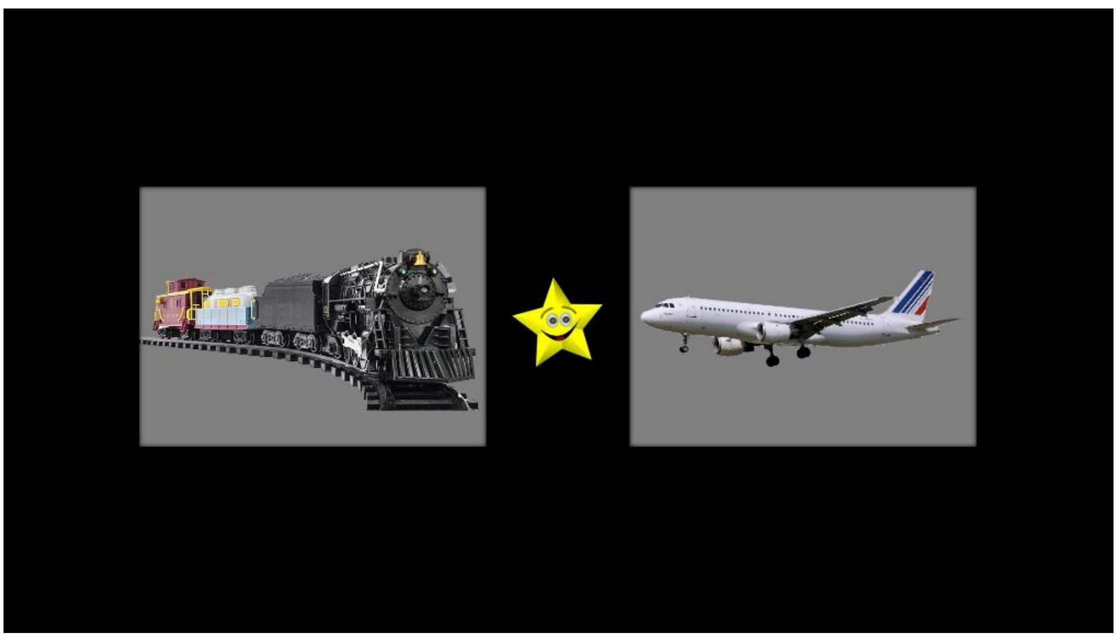

Figure 3. Example of visual stimuli in an experimental trial. Each trial began with a preview period, during which target and distractor images appeared alone, side by side. After 1500 ms a salient central image (e.g., smiling star) appeared. Simultaneously, the attention getting auditory stimulus "Look!" was presented. Once the toddler fixated on the central image for 100 ms, it disappeared. Then, the spoken label for the target was presented, followed by an encouraging phrase (e.g., "Train! You're doing great!").

Procedure

Before coming to the lab, parents were asked to complete an online version of the MacArthur Bates Communicative Developmental Inventory, Words and Sentences from (MBCDI:WS; Fenson et al., 2006). For the laboratory visit, an experimenter greeted the family at the building entrance and led them to the laboratory. After arrival of the family, one experimenter explained the procedure to parents and provided informed consent, while another experimenter played with the child to accustom them to the lab environment. Once the child was comfortable, parent and child were led to an adjacent room to complete the eye-tracking task. After the eye-tracking task, the parent and child were led back to the first room, where parents completed three questionnaires – a background history form, a vocabulary checklist for the experimental items, and the short form of the Early Childhood Behavior Questionnaire (ECBQ; Putnam, Garstein, & Rothbart, 2006) – while the child played with an experimenter.

Eye-tracking procedure and apparatus

For the eye-tracking task, the child was seated in a toddler car seat in front of a 24-inch monitor. The parent sat to the left and slightly behind the child, while an experimenter sat to the right so as to be able to encourage the child to attend to the screen. Before the experiment, parents were asked to refrain from naming or describing any of the images on the screen.

The experiment began with a short sesame street video so as to keep the child focused on the screen while the experimenter adjusted and focused the apparatus. Then, the tracker was calibrated using a five-point procedure with a looming bulls' eye image paired with a whistling sound. For the experimental task, each trial began with a small (30x30-pixel), colorful gaze-contingent central image presented on a black background. The image disappeared after the child fixated on it, and was replaced by the target and distractor images, which were displayed side by side (see Figure 1). After 2000 ms, another colorful central image appeared (100x100-pixel). Simultaneously, the attention getting auditory stimulus “Look!” was presented. Once the toddler fixated on the central image for 100 ms, it disappeared. Then, the spoken label for the target was presented, followed by an encouraging phrase (e.g., “Train! You’re doing great!”). The target and distractor images were displayed on the screen for 4000 ms post label-onset.

Each experimental item appeared four times across the experiment, counterbalanced so that it appeared twice as the target and twice as the distractor, once on each side of the screen for both. There were 24 experimental trials, distributed across three blocks of eight trials. In each block, there were also four trials containing novel objects. These trials are part of a separate study and not discussed further in this paper. Additionally, every four trials, children saw filler trials that contained scenes with popular characters. These scenes were accompanied by encouraging phrases (e.g., “Wow! Look at that!”). The entire eye-tracking procedure lasted approximately 15-20 minutes.

The movements of participants’ right eyes were recorded from image onset to offset at 500 Hz, using an SR-Research Eyelink 1000 plus eye-tracker. Movements were binned into 50 ms intervals and classified into fixations, saccades, and blinks following the default setting on the eye-tracker. Finally, target and distractor areas of interest were defined as the two 650x400-pixel areas where the corresponding images were displayed.

Offline assessments

Parents completed four offline questionnaires: the MBCDI:WS, a background history questionnaire, a vocabulary checklist for the experimental items, and ECBQ short form.

After scheduling their visit, parents were emailed a link to an online version of the MBCDI:WS, but encouraged to fill it out shortly before the experiment. If parents filled out the form more than a month before the lab visit, they were asked to fill it out again at or shortly before the visit. The remaining forms were filled out at the lab visit, after the eye-tracking task. For the background history questionnaire, parents were asked to provide demographic information and answered questions regarding their toddler's language environment. For the vocabulary checklist, parents rated their child's comprehension and production for all of the experimental items on a scale of 1 (child definitely does not understand/say the word) to 4 (child definitely understands/says the word). Items for which parents parked comprehension as less than a '3' were removed from subsequent analyses. Finally, the ECBQ asks about a range of temperamental characteristics of the child by presenting a description of a behavior and asking the parent to describe how often they have observed their child doing the behavior in the last two weeks on a scale from 1 (never) to 7 (always). For example, one question asks, "When engaged in play with his/her favorite toy, how often did your child play for more than 10 minutes?" The ECBQ measures a number of aspects of child temperament, however we limited our analyses to only include the attentional focusing subscale.

Data exclusion

As laid out in the pre-registration (Appendix C), individual trials were removed if the percentage of track-loss exceeded 80 percent, following prior precedent (e.g., Borovsky et al., 2016). Furthermore, the data for individual participants was removed if, after the removal of individual trials based on track-loss and the vocabulary questionnaire, they did not have at least two high connectivity and two low connectivity data points.

Results

Approach to analyses

Analysis of fixation data

We explored how looks to items before and after labeling relate to the connectivity characteristics of the items (normative High/Low and corrected individual degree) and the three individual differences measures (age, vocabulary percentile, and attentional focusing score). For the pre-labeling analyses, we calculated the proportion of fixations to the normative high connectivity item versus the normative low connectivity item over the full preview time period consisting of the 1,500 ms time window going from stimuli image onset to central image onset. For the post-labeling analyses, we calculated the proportion of fixations to the target versus the distractor over the pre-registered test period consisting of the 1,500 ms time window going from 300 to 1,800 ms post label onset – corresponding with time windows used in other studies of infant known-word recognition (e.g., Borovsky et al., 2016; Fernald et al., 2008).

For both sets of analyses, we compared interest areas using a dependent measure calculated by taking the log of the ratio of fixations to the relevant items: high connectivity over low connectivity in the preview period, and target over the distractor in the test period. For this metric, positive values indicate a bias towards the numerator (high connect item and target), negative scores indicate a bias towards the denominator (low connect item and distractor), and a score of zero indicates equivalent looks between the two. This log gaze-ratio measured is commonly used to circumvent model assumption violations that result when using proportions, which are bounded between 0 and 1 (e.g., Arai, van Gompel, & Scheepers, 2007; Borovsky et al., 2016; Knoeferle & Kreysa, 2012).

Planned analyses noted in the pre-registration (Appendix C) consisted of time window analyses using linear mixed effect modeling to explore how degree and the individual differences measures relate to processing in the preview and test periods. Exploratory analyses consisted of 1) time window analyses exploring different time windows in the preview period and the influence of item pair trial order, and 2) growth curve analyses aimed at uncovering nonlinear patterns in the data.

Calculating high-low corrected degree difference scores for preview period analyses

In the preview period, using each participant's individually calculated corrected degree values we then calculated difference scores between low and high items. We chose to do this in preview period because neither object has been singled out to be of interest to the participant, and thus any effect of connectivity on fixations is likely to result from a comparison between the two items.

Exploring the influence of semantic connectivity on pre-labeling attentional biases

Pre-labeling preview period time course

We start by visualizing and qualitatively analyzing the time course of fixations to the normative high and low category items in the pre-labeling preview period. We calculated the mean proportion of time spent fixating the two target areas (normative high and normative low item pictures) in 100 ms bins (Figure 4, Left). We then used the proportion values to calculate the high versus low log gaze-ratio for each time point (Figure 4, Right).

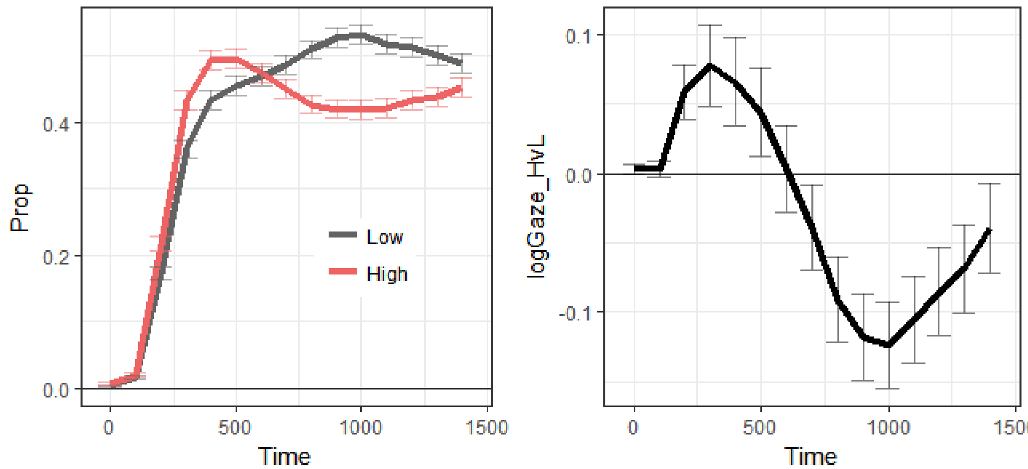

Figure 4. Preview period timecourse plots of (LEFT) the proportion of fixations to the normative high (red) and low (grey) items, and (RIGHT) the log gaze-ratio of fixations to normative high vs low items, calculated over 100 ms time bins (with *SE* bars).

In the plot of fixation proportions there are two apparent visual patterns. First, fixations to both items start near zero, as the trial only begins once they have fixated a central image between where the two items will appear, but once they begin to rise fixations to high items rise more quickly than fixations to low items. This early bias towards the high item is clearly seen as an initial rise in the log gaze-ratio plot. However, after an early peak, fixations towards the high items subside slightly and the low items become dominant. This is also clearly represented as a negative deflection in the log gaze-ratio plot.

Pre-labeling preview period time window analysis

For this analysis, we calculated the mean log gaze-ratio of looks to the normative high category item versus the normative low item over the 1,500 ms preview period going from stimuli onset to central fixation image onset (preview log gaze-ratio of High vs Low). First, a full linear mixed effects model was constructed using R (R Core Team, 2019) and lme4 (Bates, Maechler, Bolker, & Walker, 2015), with random intercepts for participant and item pair, in which the preview log gaze-ratio of High vs Low items was predicted as a function of the difference in corrected degree between the high and low items (high-low corrected degree difference), the three individual differences measures (age, percentile, and Attentional Focus score), and the three sets of interactions between high-low corrected degree difference and the individual differences measures. This full model was then fed through a backwards stepwise feature selection algorithm (set to retain random effects) using lmerTest (Kuznetsova, Brockhoff, & Christensen, 2017), which removes fixed effects based on p-values calculated at each step using Satterthwaite approximation. No fixed effects remained in the model output from the backwards selection process, indicating there were no robust predictors of fixations to high versus low items when looking at mean patterns over the full preview period.

Exploring the influence of semantic connectivity on object-label processing

Post-labeling test period time course

As with the preview period, we start by visualizing and qualitatively analyzing the time course of fixations to the target and distractor for normative high and low degree targets. We once again calculated the mean proportion of time spent fixating the two target areas (target and distractor item pictures), for normative high and low category target conditions, in 100 ms bins (Figure 5, Left). We then used the proportion values to calculate the target versus distractor log gaze-ratio for each time point, for normative high and low target conditions (Figure 5, Right).

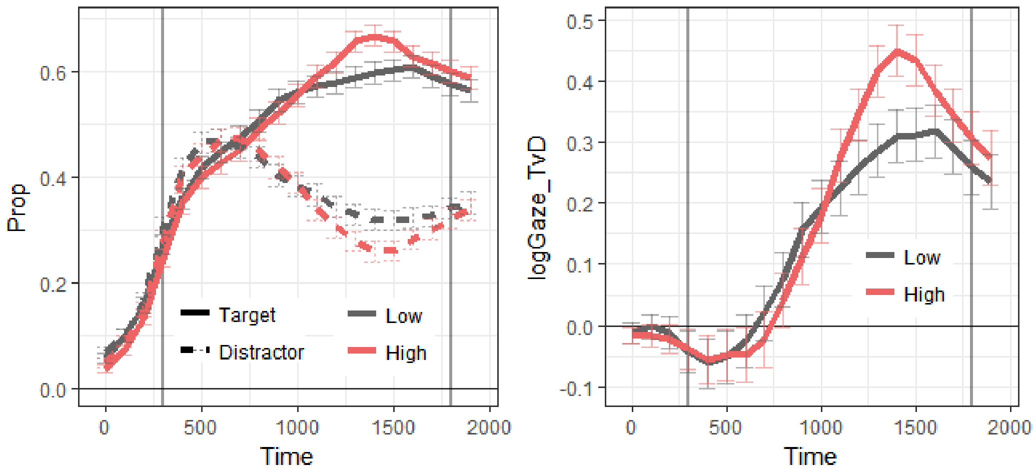

Figure 5. Test period timecourse plots of (LEFT) the proportion of fixations to the target (solid line) and distractor (dotted line) items for normative high (red) and low (grey) target conditions, and (RIGHT) the log gaze-ratio of fixations to normative high (red) vs low (grey) items, calculated over 100 ms time bins (with *SE* bars).

In the plot of fixation proportions there are two apparent visual patterns. First, fixations to both items start near zero, as the label is only presented once participants have fixated the central image between the two items, then fixations to both target and distractor begin to rise quickly until around 750 ms post label onset when fixations to distractors begin to subside while fixations to the target continue to rise. This pattern is apparent in the log gaze-ratio plot, with values hovering near zero until around 750 ms post label onset, before they begin to rise, indicating a bias towards target. Second, peaking between 1250 and 1500 ms post label onset, there is a period in which fixations towards normative high targets appear to be greater than fixations to low targets.

Post-labeling test period time window analysis

In this analysis, we calculated the mean log gaze-ratio of looks to the target versus distractor items over the 1,500 ms test period going from 300 to 1,800 ms post label onset. First, a full linear mixed effects model was constructed, with random intercepts for participant and item pair, in which the test period log gaze-ratio of Target vs Distractor items was predicted as a function of the preview period log gaze-ratio of Target vs Distractor items, normative high or low degree status, individual corrected target degree, the three individual differences measures (age, percentile, and Attentional Focusing score), and the six sets of interactions between normative status and individual corrected degree on the one hand and the individual differences measures on the other. This full model was then fed through a backwards stepwise feature selection algorithm. The coefficients for the resulting model are presented in Table 1.

Table 1. Results for the model output from the backwards stepwise feature selection of log gaze-ratio of fixations to target vs distractor in the test period.

Variable	Coef.	95% CI	p.val
Intercept	**0.156**	[0.032, 0.279]	**.03**
Preview Log Gaze-Ratio	**0.08**	[0.04, 0.122]	<**.001**
Normative High	0.052	[-0.192, 0.296]	.69
Attentional Focus	0.025	[-0.019, 0.069]	.27
Norm High: Atten Focus	**0.098**	[0.02, 0.176]	**.01**

Note. Coefficients are β_{std}.

First, the coefficient for the intercept is significantly greater than zero, quantitatively confirming that participants are recognizing and fixating the target item. Next, the coefficient for the log gaze-ratio of the target vs distractor for the preview period is also significantly greater than zero, confirming that toddler’s test period looking patterns are influenced by their initial attentional biases during the preview period – as is well documented in the Intermodal Preferential Looking Paradigm literature (e.g., Golinkoff et al., 2013). Finally, although the main effects of normative high versus low status nor attentional focus score are significant, their interaction is significantly greater than zero. This significant interaction captures the fact that while children with high Attentional Focus scores fixate normative high degree targets more than low, children with low Attentional Focus scores show the opposite pattern – as is apparent in the plot of model fit estimates of log gaze-ratio of fixations to target vs distractor items in the test period as a function of normative degree group and attentional focus scores shown in Figure 6.

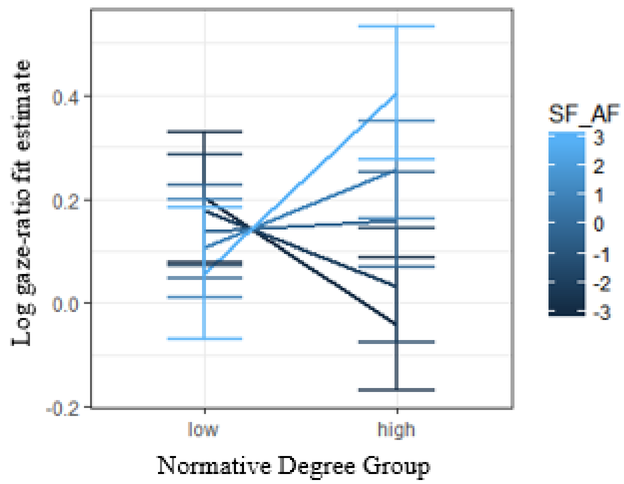

Figure 6. Model estimates of log gaze-ratio of fixations to targets vs distractor items in the target period as a function of normative degree group and attentional focusing score.

Exploratory analyses

Analyses of the influence of connectivity in exploratory time windows in the pre-labeling preview period

This and the following growth curve analysis are designed to explore the possibility that the lack of any significant predictors for mean fixations to high versus low items over the full preview period was due to the non-linear pattern that is visually apparent in the time course plots. Here we take the simple approach of dividing the preview period in half and, after exploring the relation between the two halves, following the steps of the planned analysis for each half separately.

First, we explored the relation between the first and second half of the preview period by creating a linear mixed effects model, with random intercepts for participant and item pair, in which the log gaze-ratio of High vs Low items for the second half of the preview was predicted as a function of log gaze-ratio for the first half of the preview period. Results indicated a significant negative relation, $\beta = -0.17$, $p < .001$, quantitatively confirming the inverse relation between the two halves noted in the qualitative analysis of the timecourse plots.

Next, we follow the steps of the planned analysis for each half separately. For the first half, no fixed effects remained in the model output from the backwards selection process, indicating there were no robust predictors of fixations to high versus low items when looking at mean patterns over the first half of preview period. The coefficients for the model output from the backwards selection process for the second half of the preview period are shown in Table 2.

Table 2. Results for the model output from the backwards stepwise feature selection of log gaze-ratio of fixations to high vs low items in the second half of the preview period

Variable	Coef.	95% CI	p.val
Intercept	0.002	[-0.249, 0.254]	.98
Corrected Degree Difference	**0.103**	[-0.004, 0.201]	.**04**
Percentile	0.032	[-0.035, 0.097]	.35
Corrected Deg Dif: Percentile	**0.092**	[0.016, 0.164]	.**01**

Note. Coefficients are β_{std}.

The results indicate a significant positive main effect of individual corrected degree difference between high versus low items, and a positive interaction between corrected degree difference and percentile. We visualize the interaction in Figure 7 by plotting model fit estimates of (scaled) log gaze-ratio of fixations to high vs low items in the preview period as a function of corrected degree difference and percentile.

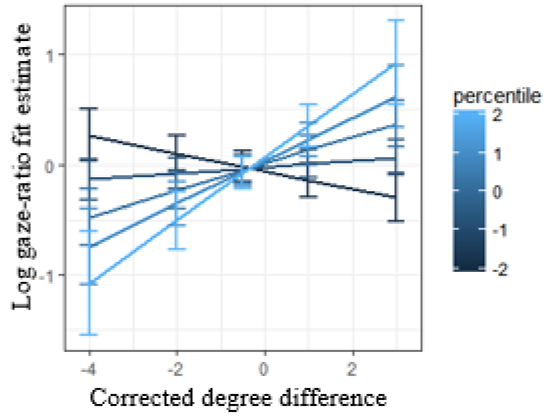

Figure 7. Model estimates of log gaze-ratio of fixations to high vs low items in the preview period as a function of corrected degree difference and percentile.

The figure clearly shows two patterns. First, we can see the significant positive effect of corrected degree difference, where log gaze-ratio fit estimates are on average higher for larger values of corrected degree difference. Second, we can see the significant interaction between corrected degree difference and percentile, with higher percentiles showing a positive relation between degree difference and fit estimate while lower percentiles showing a negative relation. In other words, in the second half of the preview period, the bias towards the normative low degree item is driven by high percentile participants paying more attention to normative low degree items for which the difference in individually calculated degree between normative high and low items is relatively smaller and by low percentile participants paying more attention to low items for which the differences are relatively large.

Growth curve analysis of the influence of connectivity on looking patterns during the pre-labeling preview-period

We performed a growth curve analysis (Mirman, Dixon, & Magnuson, 2008; Mirman 2014) to compare the timecourses of fixation proportions to normative high versus low category items using orthogonal polynomial timecodes (i.e., uncorrelated linear, quadratic and cubic components of the time predictor, which are better suited to multiple regression) calculated using eyetrackingR (Dink & Ferguson, 2018). The model consisted of a linear mixed effects model, with random intercepts and slopes for participant and item pair, in which the proportion of fixations to the item was predicted as a function of the normative high vs low status of the item, the first three orthogonal polynomial timecodes, and the interactions between high vs low status and the timecodes. The results for the analysis are presented in Table 3.

Table 3. Results for growth curve analysis of the timecourse of fixation proportions to normative high vs low items across the preview period.

Variable	Coef.	t.val	p.val
Intercept	**0.386**	128.92	**<.001**
Normative High	**-0.013**	-5.29	**<.001**
OT 1	**0.459**	45.8	**<.001**
OT 2	**-0.366**	-38.74	**<.001**
OT 3	**0.168**	16.98	**<.001**
Normative High: OT 1	**-0.085**	-9.041	**<.001**
Normative High: OT 2	0.011	1.15	.25
Normative High: OT 3	**0.075**	8.0	**<.001**

Focusing on the coefficients of interest, the main effect of normative high versus low status and the interactions of normative status with the linear and cubic polynomial timecodes significantly contributed to the model. To aid in deconstructing the meaning of these effects, in Figure 8 we plot the full model, as well as four separate component models including only the main effect and it's interactions with each of the linear, quadratic and cubic orthogonal timecodes. The negative coefficient for the main effect of normative high vs low degree status indicates that on average over the full preview period participants make fewer fixations to the normative high items. The significant negative interaction between normative high vs low and the linear timecode captures how fixations towards normative high items are higher than towards normative low items towards the beginning of the time window, but then fixations to normative low items are greater towards the end of the time window. Finally, the significant positive interaction between normative high and the cubic time code captures how the pattern of fixations towards the normative high item is more complex than for low item – with an initial peak, followed by a trough, again followed by a gentle rise – in comparison to a single broad peak for fixations to the low item.

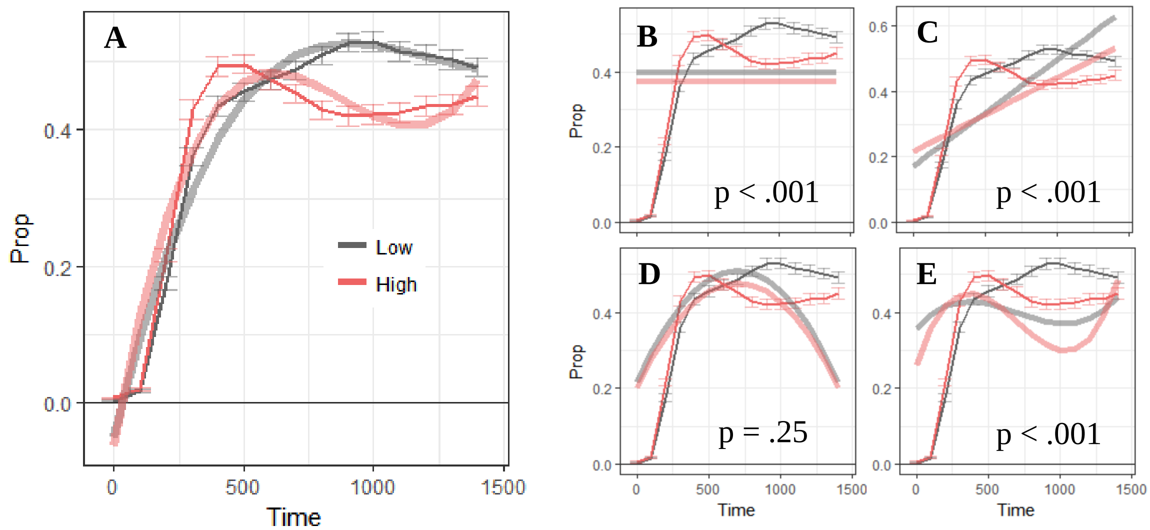

Figure 8. Preview period timecourse plots of the proportion of fixations to the normative high (red) and low (grey) items, overlaid with (A) the full growth curve model, and component models demonstrating (B) the main effect of normative high vs low, and the interactions between normative high vs low and the (C) linear, (D) quadratic, and (E) cubic orthogonal timecode polynomials, calculated over 100 ms time bins (with *SE* bars).

Analysis of relations between attentional biases in the exploratory preview period time windows and looking patterns in the post-labeling test period

To explore what portion of the preview period better predicts fixation patterns in test period, we replicated the planned analysis from section 3.3.1, but with the inclusion of preview period log gaze-ratios of Target vs Distractor items calculated over the full preview period (full preview LGR), the first 750 ms of the preview period (1st half preview LGR), and the last 750 ms of the preview period (2nd half preview LGR). This full model was then fed through a backwards stepwise feature selection algorithm. The coefficients for the resulting model are presented in Table 4.

Table 4. Results for the model output from the backwards stepwise feature selection of log gaze-ratio of fixations to target vs distractor in the test period with three versions of preview period log gaze-ratio.

Variable	Coef.	95% CI	p.val
Intercept	**0.153**	[0.032, 0.274]	.**03**
1st Half Preview LGR	**0.097**	[0.057, 0.139]	<.**001**
Normative High	0.043	[-0.195, 0.282]	.71
Attentional Focus	0.025	[-0.019, 0.07]	.26
Norm High: Atten Focus	**0.093**	[0.015, 0.17]	.**02**

Note. Coefficients are β_{std}.

The model output from the backwards selection algorithm is nearly identical to that from the planned analysis, with the main difference being that the log gaze-ratio for the first half of the preview period remains in place of that for the full preview period. This is due to the nonlinear pattern seen in Figure 8, and the fact that the initial bias towards the normative high item – versus the bias towards the normative low item in the second half of the preview period – maps onto the pattern of greater fixations for normative high compared to low targets in the test period.

Analysis of order effects in post-labeling test period

In this analysis, we explore if any of the effects described in section 3.3.1 vary over the over the course of the four presentations for each pair of items by entering item pair presentation order into the model as a main effect and as interactions with all other elements of the model. This full model was then fed through a backwards stepwise feature selection algorithm. The coefficients for the resulting model are presented in Table 5.

Table 5. Results for the model output from the backwards stepwise feature selection of log gaze-ratio of fixations to target vs distractor in the test period with item pair order included.

Variable	Coef.	95% CI	p.val
Intercept	**0.214**	[0.066, 0.362]	**.01**
Preview Log Gaze-Ratio	**0.185**	[0.087, 0.286]	**<.001**
Pair Order	-0.024	[-0.061, 0.013]	.20
Normative High	0.058	[-0.18, 0.295]	.65
Attentional Focus	0.024	[-0.02, 0.068]	.29
Preview LGR: Pair Order	**-0.04**	[-0.076, -0.006]	**.02**
Norm High: Atten Focus	**0.104**	[0.026, 0.181]	**<.01**

Note. Coefficients are β_{std}.

The output from the backwards selection algorithm is once again very similar to as in the original analysis, with the exception that pair order remains in the model due to a significant negative interaction with preview period log gaze-ratio. In other words, over the course of the four presentations of a pair of items the influence of the preview period lessens. This lessening influence of the preview period in combination with the fact that counterbalancing did not include reverse order presentations may be underlying the slight preference for the distractor seen in the time course plots in Figure 5 around 500 ms post label onset which emerges in the last two presentations of item pairs.

Discussion

Building on recent modeling work indicating early lexico-semantic development, at the group level, is related to patterns of shared perceptual features (Peters & Borovsky, 2019), this project explored whether patterns of shared perceptual features among noun-concepts influence processing at the individual level. Furthermore, we asked whether any effects of perceptual similarity structure on processing might relate to differences in patterns of pre-labeling attention. This question was based on recent findings tying the semantic structure of children's early productive vocabularies to word processing skill on the one hand (e.g., Borovsky et al., 2016) and use of attentional biases on the other (e.g., Colunga & Sims, 2017). One explanation that ties these results together is that children employ feature-guided selective attention based on patterns of shared perceptual features between to-be-processed items and known vocabulary. Such processes could influence what objects children attend to in their environments, jumpstart the creation of cohesive representations and aid in initial label learning (Jones et al., 1991), and push attended lexical items further along the novel-familiar continuum – resulting in facilitated processing (McMurray et al., 2012). Building on this previous work and using a recently developed set of semantic features norms for early acquired words (Peters et al, in prep) in combination with a variant of the Looking While Listening Paradigm (Borovsky et al., 2015; 2016), we asked whether toddlers' visual attentional biases to known objects and subsequent processing of object-labels is influenced by objects' patterns of (visual-motion and visual-form and surface) perceptual-semantic connectivity with other noun-concepts in individual toddlers' productive vocabularies. As a secondary focus, we explored whether any effects of connectivity relate to individual differences in age, vocabulary percentile, and temperamental tendency to maintain focused attention.

To explore the influence of connectivity on visual biases we compared looks to high versus low (normative and individually calculated) perceptual-semantic connectivity objects in the pre-labeling preview period of the task. At the outset, we noted three possible outcomes of interest: 1) items with more (visual-motion and visual-form and surface) featural connections with other known words might draw attention, 2) we might see greater looks to low connectivity items, or 3) we could see dynamic changes in attentional biases over the course of the preview period. Planned time window analyses of mean looking over the whole preview period indicated no robust preferences for either high or low connectivity items. However, qualitative analyses of the timecourse plots indicated attention to the normative high connectivity items is fast and first, and attention to the normative slow connectivity items is slow and second, which offers tentative support for the third proposed outcome. The qualitative analysis was supported by exploratory time window and growth curve analyses. Finally, exploratory time window analyses also uncovered individual differences in the influence of individually calculated perceptual connectivity on attentional biases in the latter half of the preview period relating to productive vocabulary percentile – indicating the bias towards the normative low degree item is driven by a combination of high percentile participants paying more attention to normative low degree items for which the difference in individually calculated degree is relatively larger and low percentile participants paying more attention to low degree items for which the differences in individually calculated degree are relatively larger.

We investigated how the connectivity of noun-concepts influences object-label processing and its relation with pre-labeling biases by comparing looks to target versus distractor items as a condition of connectivity in the post-labeling test period. In addition to exploring whether connectivity facilitates or inhibits processing, we were interested in whether or not the pattern of results in the test period would naturally follow from those of the preview period. Qualitative analyses of the timecourse plots indicated that processing is facilitated for normative high connectivity items, with the planned time window analysis indicating the effect is largely driven by participants with high Attentional Focus scores. Importantly, this effect was significant even when considering the mean pattern of fixations over both the whole preview period and the first half of the preview period – which turned out to be a better predictor of test period fixations in exploratory analyses. The fact that there is an initial bias towards high connectivity items in the preview period that parallels the facilitatory effect of label processing for those same items in the test period indicates there could be some cascading effects in play. However, the fact that the effect of high connectivity comes out even when considering the biases from the first half of the preview period indicates either the presence of an additional processing boost or that the cascading effects of initial biases are nonlinear. Indeed, one such nonlinear effect was revealed by the significant interaction between trial order and preview period looking patterns, demonstrating that cascading effects differ over the course of the experiment. Other such effects might be revealed by taking a more granular look at patterns of fixations within trials, or by including non-labeling trials in the experimental design.

Altogether, these results offer tentative support for the idea that early noun-concept processing is influenced by feature-guided selective attention related to patterns of shared (visual-motion and visual-form and surface) perceptual features between to-be-processed items and known vocabulary. Through the lens of this mechanism, the initial bias towards the high connectivity items in the preview period can be explained as resulting from the recognition of shared perceptual features drawing attention to the high connectivity items. Furthermore, the pattern of individual differences in the second half of the preview period can be explained as resulting from differences in the relative salience of the high connectivity items that results from individual differences in the number of possible shared features resulting from differences in vocabulary size. While the slower, but more robust bias towards low connectivity items was unexpected, one potential explanation is that the pattern results from the relative ease of processing high connectivity items – which would allow for greater attentional resources to be dedicated to more informative, but difficult to process low connectivity items. Moving to the test period, given the significant positive relation between attentional focus scores and percentile, the facilitatory effect of high connectivity items for toddlers with high attentional focus scores could be explained by an additional processing boost that results from a history of relatively greater attention to the high connectivity items relative to toddlers with low attentional focus scores.

We next consider some limitations that constrain the scope of these findings and point the way for future directions. First, the current project only explored the role of perceptual connectivity on attentional biases and label processing in situations in which two items from the same category were presented together. Furthermore, the item categories were limited to vehicles and animals. However, if perceptual feature-driven selective attention is indeed an important mechanism influencing early lexico-semantic development, we would expect to find effects of perceptual connectivity in a wide variety of situations, including contexts containing sets with a greater number of items from different and wider ranging categories. Furthermore, we might expect effects to extend beyond eye fixations to include differential patterns of interaction with high versus low connectivity items in more ecologically valid contexts. Second, it is difficult to determine if the results of the preview period are truly resulting from differences in connectivity or unrelated differences in relative low-level visual salience between images. While efforts were made to equate pairwise saliency of images and the use of linear mixed effects modeling with random effects for items is designed to reduce this problem, the potential that such low-level saliency effects are obscuring perceptual connectivity effects remains. Future work might benefit from either 1) making use of visual salience mapping algorithms to include metrics of relative salience as independent variables in models or 2) using some kind of more elaborate stimuli pairing design, so that some items are high connectivity in one pair and low connectivity in another. Finally, building on work demonstrating that perceptual features matter most for this age range (Peters & Borovsky, 2019), this work focused on only visual-motion and visual-form and surface features. While the results support the influence of such features in early lexico-semantic processing, future work should explore a wider range of features across a wider range of ages.

In sum, the current project provides evidence that shared (visual-motion and visual-form and surface) perceptual features between a to-be-processed noun-concept and known vocabulary items influence early lexico-semantic processing, including both pre-labeling attentional biases and subsequent processing of object-labels. This pattern of results supports accounts emphasizing the early importance of perceptual features, and provides tentative evidence for shared-feature-guided selective attention as being one of the mechanisms underlying early individual differences in online lexico-semantic processing – which could over time relate to individual differences in productive vocabulary size and language skills more broadly.

CHAPTER 3. PERCEPTUAL-SEMANTIC CONNECTIVITY INFLUENCES ATTENTION TO NOVEL OBJECTS AND LEARNING OF OBJECT-LABEL PAIRINGS

Introduction

Early differences in toddler vocabulary growth robustly predict language-related outcomes in childhood and adolescence (Dale et al., 2003; Lyytinen et al., 2005; Rescorla, 2002;2009), which in turn have cascading effects on later academic achievement, professional attainment, and social and emotional functioning (Bishop & Leonard, 2000; Bronwlie et al., 2004; Conti-Ramsden & Durkin, 2016). While these predictive links are well documented, the mechanisms underlying individual differences in early vocabulary size remain poorly understood. Recent findings that separately tie the semantic structure within individual children's early productive vocabularies to word processing skill on the one hand (e.g., Borovsky, Ellis, Evans, & Elman, 2016), and use of attentional biases to perceptual features of objects on the other (e.g., Colunga & Sims, 2017), hint at an interesting possibility. Namely, noun-concept processing may relate to the semantic connectivity of a toddler's vocabulary via attentional biases that depend on relations between the semantic features of to-be-processed words and known vocabulary. Here, we build on recent modeling work showing perceptual features matter most in early lexico-semantic development (Peters & Borovsky, 2019) and ask: Does toddlers' attention to novel objects vary with the number of the objects' perceptual-semantic relations with the words toddlers already know – or in other words, with the objects' patterns of perceptual-semantic connectivity? Furthermore, do these relations affect subsequent novel object-label learning outcomes?

A range of research provides evidence that toddlers' patterns of attention relate to their word learning outcomes. This research has explored attention at multiple levels of analysis, including: participant level measures of temperamental attentional characteristics (Dixon & Shore, 1997; Dixon & Smith, 2000), task level measures of attention in a word learning context (Kannass & Oakes, 2008), and stimuli level measures exploring relations between attention allocated to specific stimuli and subsequent learning outcomes (MacRoy-Higgins & Montemarano, 2016; Pereira, Smith, & Yu, 2014; Yu & Smith, 2011). Furthermore, a rich line of research has explored how attentional biases towards specific object features, in particular perceptual features such as shape, develop over the course of early word learning and subsequently facilitate it (e.g., Colunga & Smith, 2008; Smith, 1995, 2000; Smith, Jones, & Landau, 1996; Smith, Jones, Landau, Gershkoff-Stowe, & Samuelson, 2002). This research has revealed broad relations between word knowledge and attention – for example showing that late talkers and typically developing toddlers differ in their usage of feature biases (e.g., Colunga & Sims, 2017). However, less is known about more fine-grained effects of semantic knowledge on attention to novel objects and cascading effects on subsequent word learning outcomes.

Indirect evidence that word knowledge might influence attention to novel objects comes from recent work exploring relations between word knowledge and the learning of novel object-labels. Borovsky and colleagues (Borovksy, Ellis, Evans, & Elman, 2016) found that learning outcomes for novel object-labels was facilitated when children knew a larger proportion of the words from the category for which the novel object belonged. The degree of learning was examined using a modified version of the Looking While Listening Paradigm (e.g., Fernald et al., 2008), which characterizes online processing based on patterns of eye fixations of object images made in response to spoken labels. In the post-labeling test period, children more robustly looked to the labeled object when it was from one of their high-density categories – a result that was interpreted as facilitated label processing. However, another unexplored possibility is that the pattern of post-label looking was in part due to pre-labeling attentional biases. In other words, category density might influence attentional biases to objects, perhaps with greater looks to high-density items, which could in turn have cascading consequences on subsequent label processing. Importantly, such attentional biases could be at play throughout the entire learning process – a possibility we explore in this paper.

Relations between children's word knowledge and attentional biases to object features can be explained by the Attentional Learning Account (e.g., Smith & Colunga, 2008; Smith & Samuelson, 2006); according to which, children attune themselves to coherent structure in linguistic and nonlinguistic input that is revealed across naming situations, learning to modify their online processing to selectively attend to features that have associated with analogous situations in their past experiences (Smith et al., 2002). Thus, the facilitatory effect of high category density on novel object-label learning could be explained as resulting from the fact that children who know a higher proportion of objects within a category likely have more experience with the structure of features across items in that category. Consequently, such relatively abundant experience may result in greater attunement and sensitivity to structure in online experience, allowing children with high density categories to more easily recognize key features of novel objects and subsequently employ feature-guided selective attention based on their knowledge of the correlational structure of the category. Importantly, such feature-guided selective attention is proposed to expedite the initial creation of cohesive lexico-semantic representations (Jones et al., 1991), which could facilitate subsequent object-label mapping. Although most research in this area has focused on broadly applicable biases to feature domains such as shape, the type of shared feature-guided processes that could underly the facilitatory effect of high category density should be more idiosyncratic to the groups of objects making up such categories. Thus, to explore this issue, we need to move beyond simple descriptions of categories or feature domains, and work with more granular featural descriptions of objects to explore patterns of shared features in individual toddlers' vocabularies.

Here, we use the approach of modeling productive vocabularies using noun-feature networks. This approach builds on recent work applying graph-theoretic network modeling techniques to lexico-semantic representations (e.g., Hills et al., 2009a, 2009b; Steyvers, & Tenenbaum, 2005). Noun-feature networks consist of a set of nodes connected by links – where the nodes represent noun-concepts and the links represent shared features according to a set of feature production norms (e.g., McRae et al., 2005). According to Cree and McRae's (2003) Brain Region Knowledge Type taxonomy, such features can be categorized into one of several possible types – of which perceptual features have recently been demonstrated to be of particular importance for early word learning. Perceptual features refer to information that is available through one of the five senses (e.g., has four legs). Recent work using the noun-feature approach to model early lexico-semantic development demonstrated that perceptual features – in particular the visual-motion and visual-form and surface perceptual feature sub-types – are most robustly related to the order of early word learning (Peters & Borovsky, 2019). Building on this work, here we ask whether toddlers' attention to novel objects throughout the learning process is influenced by the objects' patterns of (visual-motion and visual-form and surface) perceptual connectivity with other noun-concepts in individual toddlers' vocabularies.

The current study

The primary goal of this project is to explore how the (visual-motion and visual-form and surface) perceptual-semantic connectivity of novel objects influences 24- to 30-month-old toddlers' patterns of attention throughout the process of learning about object features and labels, and subsequent object-label learning outcomes. A secondary goal is to determine whether any semantic connectivity effects are related to age or individual differences such as word learning skill and tendency to maintain focused attention. As discussed in the previous chapter, each of these should orthogonally constrain the correlational structure of an individual's word knowledge, and thereby influence the online use of shared-feature-guided selective attention.

To explore these issues, we designed a novel object feature and label learning study. In the first stage, the Feature Exposure Task, toddlers' were introduced to two novel objects and shown videos demonstrating each object's feature, such as <flies> or <climbs trees>. For each feature pair, one feature was common in the productive vocabularies of children in the target age-range and one feature was relatively rare (according to a set of feature production norms for early learned nouns; Peters, McRae, & Borovsky, in prep). This stage includes two trials in which the novel objects were shown side by side – the first shown before the feature videos, designed to measure initial visual biases, and the second shown after the videos, designed to measure changes in biases resulting from learning about the object features. In the second stage, the Naming Exposure Task, each object was given a novel label. Finally, in the third phase, the Retention and Recognition Task, toddlers' were tested on their retention of the object labels via the aforementioned variant of the Looking While Listening Paradigm employed by Borovsky and colleagues (2016). Importantly, in this third phase, we measure both pre- and post-labeling attention, so that in addition to measuring label retention, we can explore whether attentional biases to object pairs differ after being assigned labels. Where attentional biases and label processing, in trials where object pairs are presented side by side, are characterized based on the proportion of looking time to each item in the pair, and attention during learning opportunities is characterized as the proportion of looking time out of the length of the trial. This design allows us to explore changes in attentional biases throughout the course of learning, the degree to which children attend to the feature and label learning opportunities, and whether these factors relate to eventual novel-label learning outcomes.

Simply put, the central hypotheses driving this project are that the semantic connectivity of novel objects will influence (1) toddlers' patterns of attention (during learning and subsequent free-viewing opportunities) and (2) degree of subsequent object-label retention. We consider four possible outcomes, each of which have different implication for mechanisms underlying how attention and connectivity support learning:

1. Attentional biases drive looks to novel objects with high connectivity features during learning and free-viewing opportunities, and support subsequent retention of object-label associations for those objects.
2. Attentional biases drive looks to novel objects with low connectivity features during learning and free-viewing opportunities, and support subsequent retention of object-label associations for those objects.
3. High semantic feature connectivity supports the retention of object-label associations, irrespective of attentional biases during preceding learning and free-viewing opportunities.

4. Low semantic feature connectivity supports the retention of object-label associations, irrespective of attentional biases during preceding learning and free-viewing opportunities.

As for known words in chapter 2, outcome 1 is the predicted outcome as it follows naturally from the above explanation for the facilitative effects of high category density on novel object-label learning based on the Attentional Learning Account. In contrast, outcome 2 could result if distinctiveness or novelty – which could both be defined as low similarity-based connectivity – are primary drivers of attention early in the learning process. Both of these first two outcomes – in which patterns in the retention test period parallel previous patterns of attentional biases – would lend support to the hypothesis that connectivity driven attentional biases towards objects in the environment influence object-label learning. In contrast, outcomes 3 and 4 – in which processing patterns in the retention test are unrelated to preceding attentional biases during learning and free-viewing opportunities – would indicate that semantic connectivity influences attentional biases and label learning in different ways that outweigh any potential cascading effects from the former to the latter.

The above outcomes all focus on relations between connectivity and average patterns of attention over certain time windows. However, based on the results presented in Chapter 2, it is worth noting that we might see dynamic patterns within these time windows that could be masked by such time window analyses. For example, as for known words, in the pre-labeling preview period in the retention task, we might expect to see an early bias towards high connectivity items and later, but more robust bias, towards low connectivity items. Thus, we also include exploratory analyses investigating whether there are dynamic patterns of looking within time windows that relate to semantic connectivity.

Methods

Participants

Recruitment of sixty 24- to 30-month-old toddlers was carried out via community fliers and outreach events, 49 of which were included in the final analyses (24 F, 25 M). This age range was selected because (1) it covers a time of particularly dramatic change in children's lexico-semantic systems, and (2) the productive vocabularies of children in this range are large enough to enable the selection of stimuli that are appropriately balanced for patterns of connectivity. The remaining 11 children completed the study but were excluded due to: history of chronic ear infections and hearing loss concerns (n=2), neurological impairment (n=1), insufficient data due to track-loss/fussiness (n=8). All children included in the final analyses met the following pre-registered (see Appendix D) inclusionary criteria: reported normal vision and hearing, no history of chronic ear-infections, no history of neurological or cognitive impairment, not born preterm (<37 weeks) and low birth weight (<5lb8oz), not exposed to language other than English for more than 8 hours a week (~10% of waking hours).

Approach to calculating metrics of semantic connectivity

Features

In this project we used a recently developed set of feature production norms for early-learned nouns (Peters, McRae, & Borovsky, in prep), which includes all 359 noun concepts on the MBCDI:WS. There are four broad type of features in the set: encyclopedic, functional, perceptual, and taxonomic. Perceptual features have seven subtypes: smell, sound, tactile, taste, visual-color, visual-form and surface, and visual-motion. Based on recent modeling evidence

(Peters & Borovsky, 2019), here we use only two types of perceptual features – visual-form and surface and visual-motion features.

Network construction

The productive vocabularies of each individual, as measured by MBCDI:WS , were used to construct graph-theoretic lexico-semantic networks. In the networks, nodes represent nouns the individuals produce and links between nodes represent shared perceptual (specifically: visual-form and surface and visual-motion) features. Following prior precedent (e.g., Engelthaler & Hills, 2017; Hills et al., 2009a), the links were treated as unweighted, undirected edges.

For use in stimuli selection, we also constructed a normative 26-month-old's network containing all nouns with an Age of Acquisition (AoA) less than or equal to 26 months. Using the method outlined in Braginksy, Yurovsky, Marchman and Frank (2016), AoA was calculated for all of the concrete nouns on the MBCDI:WS, using 5450 administrations of the American English version of the MBCDI:WS found on the Wordbank database (Frank et al., 2017).

Calculating metrics of semantic connectivity

As with the study in chapter 2, this project uses a set of measures based around degree, one of the simplest network metrics. The *degree* of a word is defined as the number of links it has to other words in the network. Given links are defined using shared features, degree characterizes the pattern of shared features between a word and its neighbors and is therefore closely related to the research questions driving this work. The three metrics we use are: normative degree category, corrected degree, and corrected degree difference. We calculated the *normative degree* of each item as its degree in the normative 26-month-old's (visual-form and surface and visual-motion) perceptual network described above. As described in detail in the next section, paired items were rank ordered in terms of normative degree, and assigned to either "high" or "low" *normative degree categories*. For each participant, we also used the networks created using their individual productive noun vocabularies, to calculate the *corrected degree* (a.k.a. normalized degree) for each item, defined as the degree of the item divided by maximum possible degree for that individual's network (i.e., the number of words in the network minus one). We used corrected degree instead of degree because the degree of items in an individual's network are closely related to the number of items in that network, which in turn is closely related to the age and MBCDI:WS percentile of the individual. Finally, for each pair of items we calculated the *corrected degree difference* between normative high and low items as the corrected degree of the high item minus the corrected degree of the low item.

Materials

Item creation

Nonsense objects were created by combining two or three basic shapes (e.g., circle, triangle, etc.) into a single object, which was then given a simple texture and color (see Appendix B for details). The objects were created to be matched for visual complexity and saliency within yoked pairs, verified by laboratory members. Two-syllable object names were selected from the Novel Object and Unusual Name Database (Horst & Hout, 2014), and were controlled for number of phonemes and phonemic neighborhood density (see Table 17, in Appendix B for item pairings).

Three pairs of features were selected from the extended set of McRae feature norms (McRae et al., 2005; Peters, et al., in prep). Features were chosen from a subset of features that are likely to be known by toddlers, defined as having a toddler accessibility rating greater than or equal to the median value in the database. To partially control for the potential impact of correlated features, only features for animate objects were selected, and within each pair, features were matched for feature type. Pairs were chosen to have one "high connectivity" feature, found in a relatively high number of words typically found in the productive vocabularies of 24- and 30-month-olds, and one "low connectivity" feature, found in a relatively low number of words. Across lists, features were counterbalanced, however objects and names were organized into balanced yoked pairs.

Auditory stimuli

All auditory experimental stimuli and additional encouraging phrases were spoken in an infant-directed voice by a female Standard American English speaker, and recorded on a mono channel at 44.1 kHz sampling rate. Stimuli were adjusted to a mean intensity of 70 dB and mean durations of FET: 7.98 s, NET: 13.50 s, and RRT: 1.02 s (see definitions of subtasks in 2.3.1. Eye-tracking experiment structure). Encouraging phrases, such as "Isn't that cool!" and "Do you like it?", were included in the FET and RRT portions of the experiment to maintain interest. During the RRT, an attention getter, "Look!", was played at the onset of a gaze-contingent center stimulus before the onset of spoken experimental stimuli.

Visual stimuli

There were two types of visual stimuli presented in the experiment: 1) centered videos demonstrating object features or names and 2) test trials in which two objects were presented side by side. Stimuli were presented on a 1920x1080-pixel screen. Videos consisted of 1200x750-pixel animations created using the open-source 2D animation software Synfig. Videos demonstrating object features were created so as to be matched across pairs regarding scene background, video content besides the novel object (e.g., the same tree was present for videos presenting the features <climbs trees> and <flies>), and object salience (e.g., amount of movement and size). In videos presenting object names the objects were presented on a gray background and moved side to side. In the test trials, the novel object pairs were presented on the left and right sides of the screen (counterbalanced across trials), each on a 600x450-pixel gray background. See Figure 9 for an example of the visual stimuli presented in test trials. Across lists, presentation order and side were counterbalanced for all trial types.

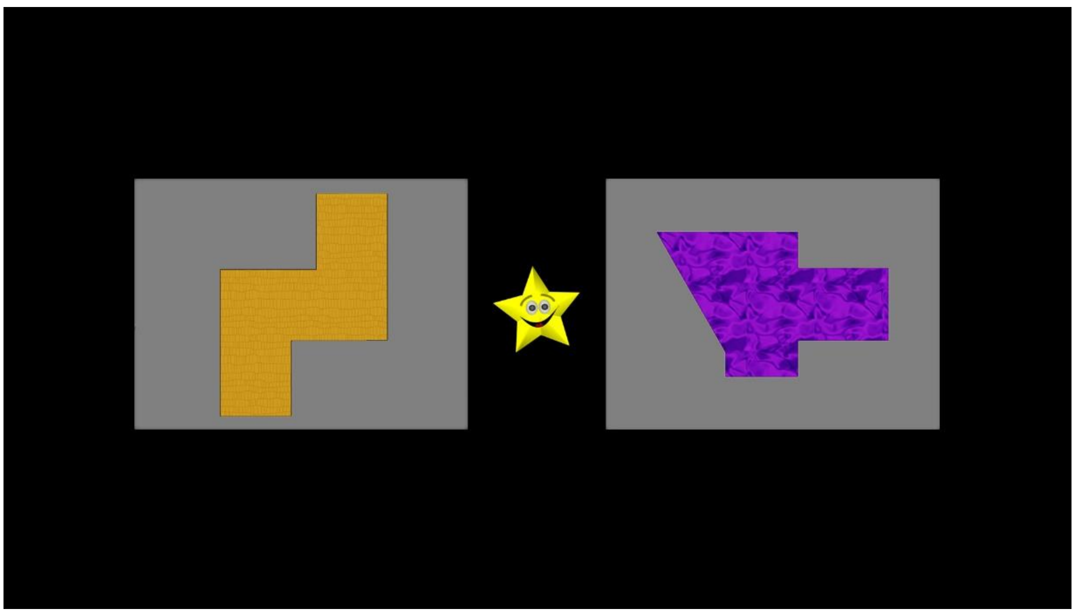

Figure 9. Example of visual stimuli in a test trial. Each trial began with a preview period, during which target and distractor images appeared alone, side by side. After 2000 ms a salient central image (e.g., smiling star) appeared. Simultaneously, the attention getting auditory stimulus “Look!” was presented. Once the toddler fixated on the central image for 100 ms, it disappeared. Then, the spoken label for the target was presented, followed by an encouraging phrase (e.g., “Train! You’re doing great!”).

Procedure

Before coming to the lab, parents were asked to complete an online version of the MacArthur Bates Communicative Developmental Inventory, Words and Sentences from

(MBCDI:WS; Fenson et al., 2006). For the laboratory visit, an experimenter greeted the family at the building entrance and led them to the laboratory. After arrival of the family, one experimenter explained the procedure to parents and provided informed consent, while another experimenter played with the child to accustom them to the lab environment. Once the child was comfortable, parent and child were led to an adjacent room to complete the eye-tracking task. After the eye-tracking task, the parent and child were led back to the first room, where parents completed three questionnaires – a background history form, a vocabulary checklist for the experimental items, and the short form of the Early Childhood Behavior Questionnaire (ECBQ; Putnam, Garstein, & Rothbart, 2006) – while the child played with an experimenter.

Eye-tracking experiment structure

The eye-tracking task consisted of three blocks, each divided into the Feature Exposure Task (FET), the Naming Exposure Task (NET), and the Retention and Recognition Task (RRT) (Figure 10). In the feature exposure task participants were first presented two novel objects side by side for 6,000 ms to determine any initial visual bias (Visual Bias). Next, they saw videos demonstrating semantic features of the novel objects (Feature Exposure; 8,000 ms). There was one video for each object, each presented twice. Finally, they were once again presented with the two novel objects side by side to determine any bias that resulted from the feature exposure (Feature Bias). In the naming exposure task participants listened to each novel object being labeled with a novel name (two trials for each object; 13,000 ms). In the retention and recognition task participants saw a pair of objects side by side, one of which was labeled with an audio stimulus. There were four trials testing retention of the novel objects (RET), and eight filler trials containing pairs of known objects. The RET trials had four time periods (determined via pilot testing): (1) Post-naming Bias Period: first 1500 ms of the trial (pre-audio stimulus), (2) Test Period: from 300 to 4000 ms post audio stimulus onset, (3) Early Test Period: from 300 to 1300 ms post audio stimulus onset, and (4) Late Test Period: from 1300 to 4000 ms post audio stimulus onset.

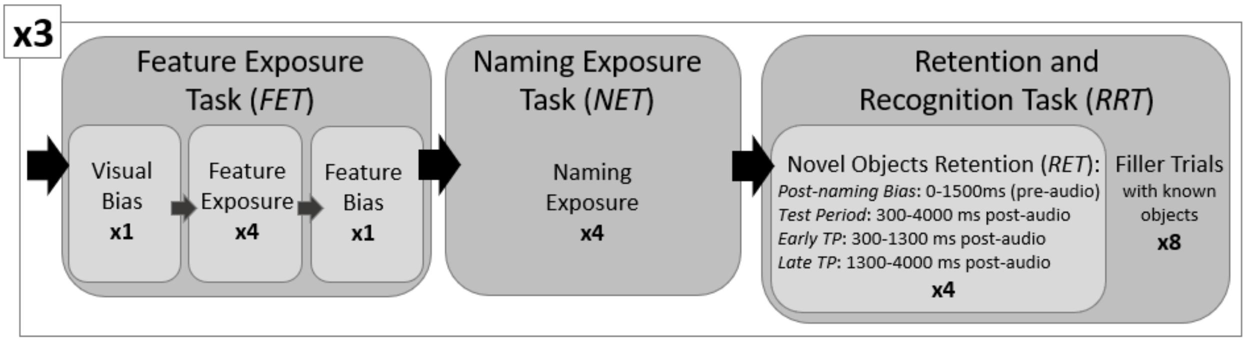

Figure 10. Diagram of one block of the experiment.

Eye-tracking procedure and apparatus

For the eye-tracking task, the child was seated in a toddler car seat in front of a 24-inch monitor. The parent sat to the left and slightly behind the child, while an experimenter sat to the right so as to be able to encourage the child to attend to the screen. Before the experiment, parents were asked to refrain from naming or describing any of the images on the screen.

The experiment began with a short sesame street video so as to keep the child focused on the screen while the experimenter adjusted and focused the apparatus. Then, the tracker was calibrated using a five-point procedure with a looming bulls' eye image paired with a whistling sound. For all tasks, each trial began with a small (30x30-pixel), colorful gaze-contingent central image presented on a black background. The image disappeared after the child fixated on it, and was replaced by visual task stimuli.

Offline assessments

Parents completed four offline questionnaires: the MBCDI:WS, a background history questionnaire, a vocabulary checklist for the experimental items, and the short form of the ECBQ.

After scheduling their visit, parents were emailed a link to an online version of the MBCDI:WS, but encouraged to fill it out shortly before the experiment. If parents filled out the form more than a month before the lab visit, they were asked to fill it out again at or shortly before the visit. The remaining forms were filled out at the lab visit, after the eye-tracking task. For the background history questionnaire, parents were asked to provide demographic information and answered questions regarding their toddler's language environment. For the vocabulary checklist, parents rated their child's comprehension and production for all of the experimental items on a scale of 1 (child definitely does not understand/say the word) to 4 (child definitely understands/says the word). Items for which parents parked comprehension as less than a '3' were removed from subsequent analyses. Finally, the ECBQ asks about a range of temperamental characteristics of the child by presenting a description of a behavior and asking the parent to describe how often they have observed their child doing the behavior in the last two weeks on a scale from 1 (never) to 7 (always). For example, one question asks, "When engaged in play with his/her favorite toy, how often did your child play for more than 10 minutes?" The ECBQ measures a number of aspects of child temperament, however we limited our analyses to only include the attentional focusing subscale.

Data exclusion

For all trials in which item pairs were presented side by side (i.e., Visual Bias, Feature Bias, and Retention Tasks), individual trials were removed if the percentage of track-loss exceeded 80 percent, following prior precedent (e.g., Borovsky et al., 2016) and as noted in the pre-registration (Appendix D). However, no track-loss cutoff was implemented for trials containing learning opportunities (i.e., feature and naming exposure videos), because high track-loss should relate to low attention, and thus provide data points that are relevant to the research questions. Data for individual participants was removed if, after the removal of individual trials based on track-loss and the vocabulary questionnaire, they did not have at least two high connectivity and two low connectivity data points in the Retention Task.

Results

Approach to analyses

Analysis of fixation data

We explored how looks to items both throughout the learning process and during test relate to the connectivity characteristics of the items (normative High/Low, degree, and clustering coefficient) and the three individual differences measures (age, vocabulary percentile, and attentional focus score).

For the Feature Exposure and Naming Exposure video trials, the dependent variable was proportion looking time, which was calculated as the ratio of the number of samples fixating the video interest area over the number of samples recorded in the trial. For the three bias (Visual, Feature, and Post-naming bias) and test period analyses, the dependent variable was calculated by taking the log of the ratio of fixation proportions to the relevant items: normative high connectivity over low connectivity for the bias period analyses, and target over the distractor for the test period analyses. For the Visual and Feature bias, we calculated the proportion of fixations over the full 6,000 ms of the trials. For the post-naming bias, we calculated the proportion of fixations over the full preview time period consisting of the 1,500 ms time window going from stimuli onset to central image onset in the retention task. Finally, for the post-labeling analyses in the retention task we calculated proportion of fixations over two pre-registered (see Appendix D) time windows determined via pilot analyses: 1) the early test period consisting of the 1,000 ms time window going from 300 to 1,300ms post label onset and 2) the late test period consisting of the 2,700 ms time period going from 1,300 ms to 4,000 ms post label onset.

For feature bias, post-naming bias, and test period analyses, we begin each section by visualizing and qualitatively analyzing the time course of fixations. We then conduct planned time window analyses. Finally, for these sections and for the feature and naming exposure sections, we also conduct exploratory analyses. Planned analyses noted in the pre-registration are presented first, followed by exploratory analyses.

Calculating high-low corrected degree difference scores

For the feature and post-naming bias analyses, using each participant's corrected degree values, we calculated difference scores between normative low and high items. We did this for these analyses because neither object has been singled out to be of interest to the participant, and thus any effect of individual connectivity on fixations is arguably likely to result from a comparison between the two items.

Exploring the influence of connectivity on pre-labeling attentional biases in the retention task

Time course

Once again, we start by visualizing and qualitatively analyzing the time course of fixations to the normative high and low items in the pre-labeling preview period. We calculated the mean proportion of time spent fixating the two target areas (normative high and normative low item pictures) in 100 ms bins (Figure 11, Left). We then used the proportion values to calculate the high versus low log gaze-ratio for each time point (Figure 11, Right).

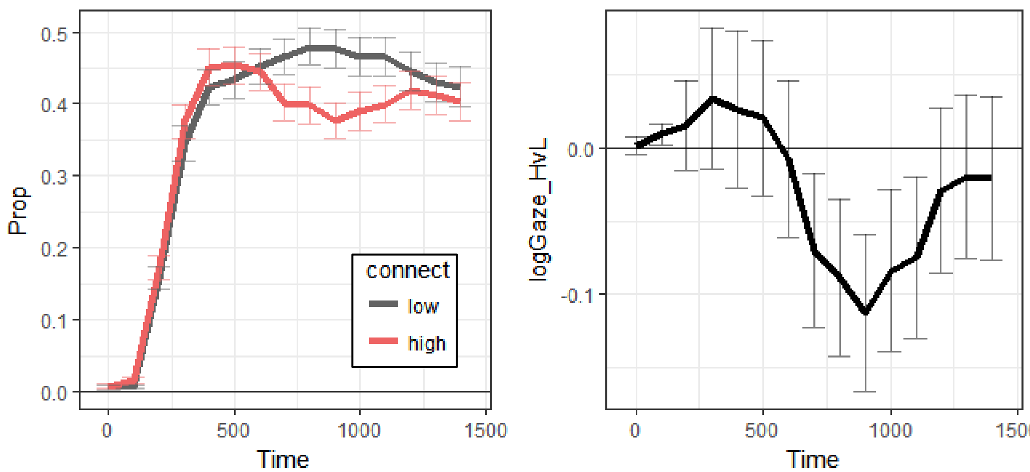

Figure 11. Preview period timecourse plots of (LEFT) the proportion of fixations to the normative high (red) and low (grey) items, and (RIGHT) the log gaze-ratio of fixations to normative high vs low items, calculated over 100 ms time bins (with *SE* bars).

In the plot of fixation proportions there are two apparent visual patterns. First, fixations to both items start near zero, as the trial only begins once they have fixated a central image between where the two items will appear, but once they begin to rise fixations to high items rise more quickly than fixations to low items. This early bias towards the high item is clearly seen as an initial rise in the log gaze-ratio plot. However, after an early peak, fixations towards the high item subside slightly and the low item becomes dominant. This is also clearly represented as a negative deflection in the log gaze-ratio plot. Remarkably, this pattern of biases closely resembles the pattern of pre-labeling biases seen for known words (this dissertation, Chapter 2).

Time window analysis

For this analysis, we calculated the mean log gaze-ratio of looks to the normative high item versus the normative low item over the 1,500 ms preview period going from stimuli onset to central fixation image onset (preview log gaze-ratio of High vs Low). First, a full linear mixed effects model was constructed, with random intercepts for participant and item pair, in which the preview log gaze-ratio of High vs Low items was predicted as a function of the Visual Bias log gaze-ratio, the difference in corrected degree between the high and low items (high-low corrected degree difference), the three individual differences measures (age, percentile, and Attentional Focus score), and the three sets of interactions between high-low corrected degree difference and the individual differences measures. This full model was then fed through a backwards stepwise feature selection algorithm (set to retain random effects). The coefficients for the model output from the backwards selection process are presented in Table 6.

Table 6. Results for the model output from the backwards stepwise feature selection of log gaze-ratio of fixations to high vs low items in the preview period

Variable	Coef.	95% CI	p.val
Intercept	-0.002	[-0.323, 0.311]	.99
Visual Bias	**0.438**	[0.085, 0.797]	**.01**
Corrected Degree Difference	0.002	[-0.116, 0.121]	.97
Age	-0.046	[-0.151, 0.059]	.39
Corrected Deg Dif: Age	**0.147**	[0.044, 0.249]	**.004**

Note. Coefficients are β_{std}.

The results indicate a significant positive main effect of Visual Bias, and a positive interaction between corrected degree difference and age. We visualize the interaction in Figure 12 by plotting model fit estimates of (scaled) log gaze-ratio of fixations to high vs low items in the preview period as a function of corrected degree difference and age.

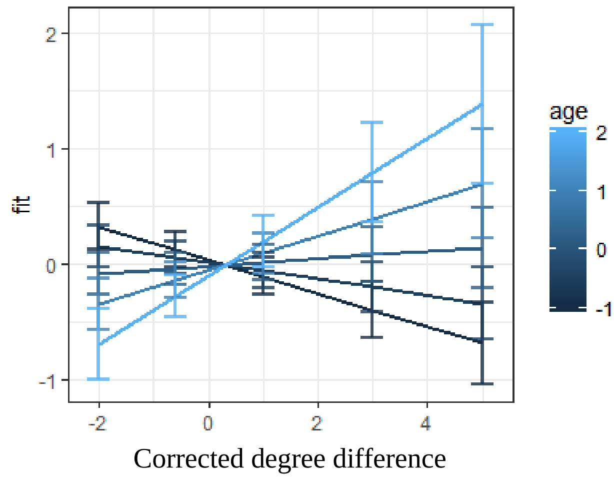

Figure 12. Model estimates (fit) of log gaze-ratio of fixations to high vs low items in the preview period as a function of corrected degree difference and age.

The positive interaction between corrected degree difference and age is clearly visible in Figure 12, which indicates that younger children are biased towards low connect items when there is a large individually calculated difference between normative high and low items, whereas the reverse is true for older children. Taken altogether, the results indicate in the preview period, after taking into account initial visual biases, the later, longer-lasting bias towards the normative low degree item is driven by a combination of younger participants paying more attention to normative low degree items for which the difference in individually calculated degree between normative high and low items is relatively larger and older participants paying attention to low degree items in cases where the differences is small.

Exploring influence of connectivity on novel object-label processing in the retention task

Time course

We first visualize and qualitatively analyze the time course of fixations to the target and distractor for normative high and low degree targets in the test period. As for the preview period, we calculated the mean proportion of time spent fixating the two target areas (target and distractor item pictures), for normative high and low target conditions, in 100 ms bins (Figure 13, Left). We then used the proportion values to calculate the target versus distractor log gaze-ratio for each time point, for normative high and low target conditions (Figure 13, Right).

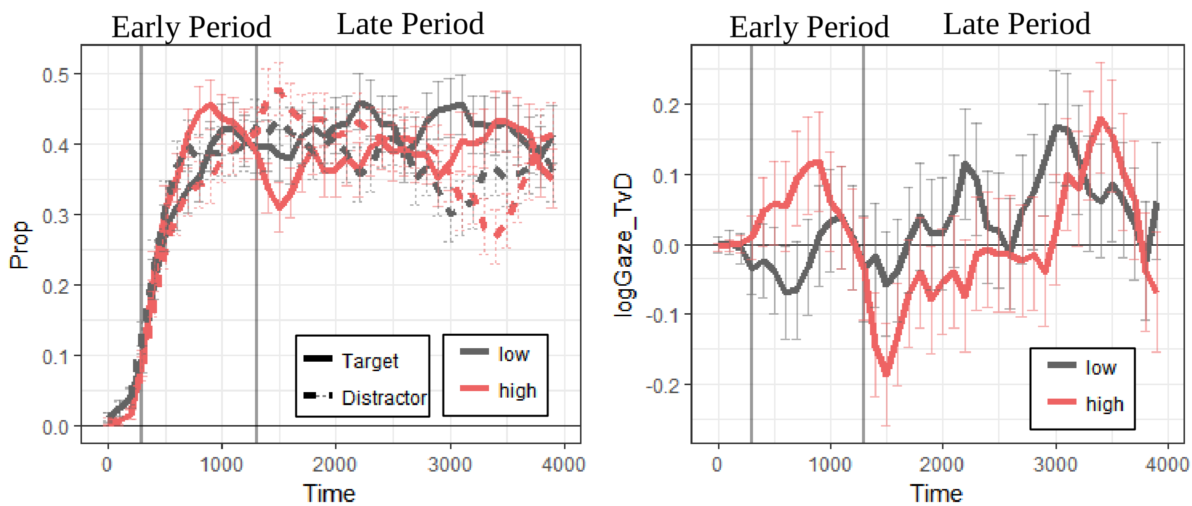

Figure 13. Test period timecourse plots of (LEFT) the proportion of fixations to the target (solid line) and distractor (dotted line) items for normative high (red) and low (grey) target conditions, and (RIGHT) the log gaze-ratio of fixations to normative high (blue) vs low (pink) items, calculated over 100 ms time bins (with *SE* bars). Early test period from 300 to 1,300 ms (marked by gray vertical lines), and late test period from 1,300 to 4,000 ms.

In the plot of fixation proportion there are three apparent visual patterns. First, as in the preview period, fixations to all items start near zero and then begin rise, whereupon we see more looks to the target than distractor in the early test period for normative high degree items but no clear difference for low degree items – more clearly visible in the log gaze-ratio plots as a rise above zero for high degree items compared to a slight dip below zero for low degree items. Second, at the start of the late test period, the pattern reverse, with greater looks to the target for normative low degree items, and greater looks to the distractor for high degree items. Finally, towards the end of the late test period, we see greater looks to the target for both normative high and low degree items. Interestingly, the broad pattern of results – early facilitation for high degree items followed by slower facilitation for low degree items – parallels the pattern of results for the preview period, in which we saw a slight initial bias towards high degree items followed by a slower but more robust bias towards low degree items.

Time window analyses

In this analysis, we calculated the mean log gaze-ratio of looks to the target versus distractor items over two time-windows: 1) the 1,000 ms early test period going from 300 to 1,300 ms post label onset, and 2) the 2,700 ms late test period going from 1,300 to 4,000 ms. First, full linear mixed effects models were constructed, with random intercepts for participant and item pair, in which the test period log gaze-ratio of Target vs Distractor items was predicted as a function of preview period log gaze-ratio of Target vs Distractor items, normative high or low degree status, individual corrected target degree, the three individual differences measures (age, percentile, and Attentional Focusing score), and the six sets of interactions between normative status and individual corrected degree on the one hand and the individual differences measures on the other. These full models were then fed through a backwards stepwise feature selection algorithm. The coefficients for the resulting models are presented in Table 7.

Table 7. Results for the model of log gaze-ratio of fixations to target vs distractor items in the early and late test period time windows

Time Window	Variable	Coef.	95% CI	p.val
Early	Intercept	-0.001	[-0.102, 0.099]	.98
	Preview log gaze-ratio	**0.115**	[0.015, 0.216]	**.03**
Late	Intercept	-0.002	[-0.115, 0.111]	.97
	Preview log gaze-ratio	**0.234**	[0.129, 0.338]	**<.001**

Note. Coefficients are β_{std}.

The results indicate that for both early and late time windows, only the preview period log gaze-ratio significantly predicts looking to target versus distractor. However, it is worth noting that the effect is stronger in the late period than in the early period.

Exploratory analyses

Analysis of influence of connectivity on amount of looking to feature exposure videos

We explored how looks to feature exposure videos related to the connectivity characteristics of the features (normative High/Low, corrected individual degree), initial visual biases, order of video presentation, and the three individual differences of interest (age, vocabulary percentile, and attentional focus score). For this analysis, we calculated the proportion of looks over the 8,000 ms feature exposure period, and constructed a full linear mixed effects model constructed using R (R Core Team, 2019) and lme4 (Bates, Maechler, Bolker, & Walker, 2015), with random intercepts for participants and feature, in which proportion of looks was predicted as a function of the visual bias (the log gaze-ratio of the proportion of looks to the item over the competitor), the order of video presentation (first or second presentation for a given feature), the two connectivity measures, the three individual differences measures, and the six sets of interactions between the connectivity and individual differences measures. This full model was then fed through a backwards stepwise feature selection algorithm (set to retain random effects) using lmerTest (Kuznetsova, Brockhoff, & Christensen, 2017), which removes fixed effects based on p-values calculated at each step using Satterthwaite approximation. However, the resulting model contained fixed effects for normative high connect and corrected degree that were highly correlated ($r = -.889$). Thus, separate models were made for each of the two connectivity measures, and separately fed through the backwards feature selection algorithm. The model containing corrected degree did not converge, while the output coefficients for the model with normative high connect are presented in Table 8.

Table 8. Results for the model output from the backwards stepwise feature selection predicting fixation proportions to feature exposure videos.

Variable	Coef.	95% CI	p.val
Intercept	-0.002	[-0.191, 0.187]	.98
Order	**-0.083**	[-0.152, -0.013]	**.02**
Visual Bias	**-0.073**	[-0.143, -0.002]	**.04**
Normative High	0.03	[-0.111, 0.171]	.68
Attentional Focus Score	**0.389**	[0.2, 0.577]	**<.001**
Norm High: AF Score	**-0.158**	[-0.299, -0.016]	**.03**

Note. Coefficients are β_{std}.

The results indicate significant negative effects of order and visual bias, a positive effect of attentional focus score, and a negative interaction of normative high degree and attentional focus score. We visualize the interaction in Figure 14 by plotting fit estimates of (scaled and centered) fixations proportions to feature exposure videos as a function of normative connect and Attentional Focus Scores, calculated using the effects library for R (Fox, 2003).

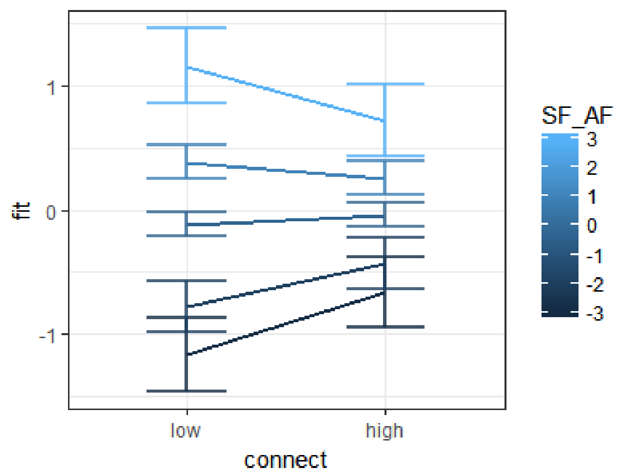

Figure 14. Model estimates (fit) of fixation proportions to feature exposure videos as a function low/high normative connect and Attentional Focus scores.

Both the positive main effect of attentional focus scores and the negative interaction with normative connect are clearly visible in Figure 14, which indicates that children with low Attentional Focus scores attend high connect items more than low, whereas the reverse is true for children with high Attentional Focus scores. Taken altogether, the results indicate that feature exposure videos are fixated more the first time than the second, more when there was an initial visual bias towards the *competitor* item, and more for participants with higher attentional focus scores – particularly for items with low normative degree.

Exploring attentional biases after feature exposure videos

Time course

We start by visualizing and qualitatively analyzing the time course of fixations to the normative high and low items over the full 6,000 ms length of the feature bias trial. We calculated the mean proportion of time spent fixating the two target areas (normative high and normative low item pictures) in 100 ms bins (Figure 15, Left). We then used the proportion values to calculate the high versus low log gaze-ratio for each time point (Figure 15, Right).

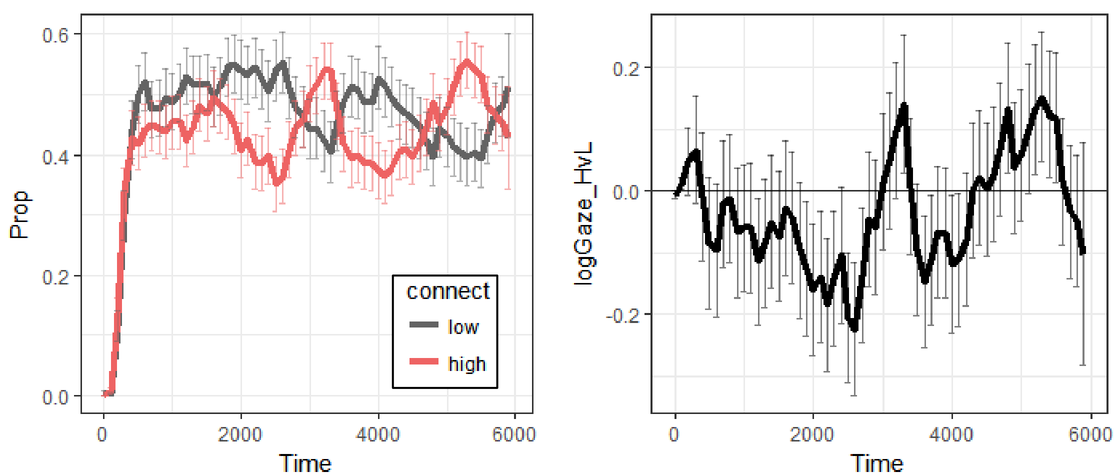

Figure 15. Feature Bias timecourse plots of (LEFT) the proportion of fixations to the normative high (red) and low (grey) items, and (RIGHT) the log gaze-ratio of fixation to normative high vs low items calculated over 100 ms time bins (with *SE* bars).

In the plot of fixation proportions there are two apparent visual patterns. First, fixations to both items start near zero, as the trial only begins once they have fixated a central image between where the two items will appear, but once they begin to rise fixations to low items rise slightly higher than fixations to high items. This early bias to low items is clearly seen as an initial drop in the log gaze-ratio plot below zero which lasts until around 3,000 ms. However, after this initial difference, from 3,000 ms to the end of the trial there is no systematic difference in fixations to high vs low items.

Time window analyses of influence of connectivity on attentional biases after feature exposure

For this analysis, we calculated the mean log gaze-ratio of looks to the normative high item versus the normative low item over the full 6,000 ms of the Feature Bias trial starting from stimuli onset (Feature Bias log gaze-ratio of High vs Low). We then built a full linear mixed effects model, with random intercepts for participant and item pair, in which the log gaze-ratio was predicted as a function of the Visual Bias log gaze-ratio, the difference in corrected degree between the high and low items (high-low corrected degree difference), the three individual differences measures (age, percentile, and Attentional Focus score), and the three sets of interactions between high-low corrected degree difference and the individual differences measures. This full model was then fed through a backwards stepwise feature selection algorithm (set to retain random effects). No fixed effects remained in the model output from the backwards selection process, indicating there were no robust predictors of fixations to high versus low items when looking at mean patterns over the time window. However, visual inspection of the relation between corrected degree difference and visual bias log gaze-ratio hinted that there may be an effect falling just outside of significance (see Figure 16). We investigate this possibility next.

Once again using Feature Bias log gaze-ratio of High vs Low, we started with a simple multivariate linear model of the Feature Bias log gaze-ratio predicted by the (similarly calculated) Visual Bias log gaze-ratio (of fixations to the objects before feature exposure) and the difference in corrected individually calculated degree between high and low items (high-low corrected degree difference). The model was significant ($F(2, 101) = 4.275$, $p = .02$, Adjusted R^2

= .06), but with only high-low corrected degree difference significantly contributing to the model (β_0 = -0.005, p = .96; $\beta_{\text{Vis Bias}}$ = 0.12, p = .20; $\beta_{\text{deg dif}}$ = -0.24, p = .01; standardized coefficients).

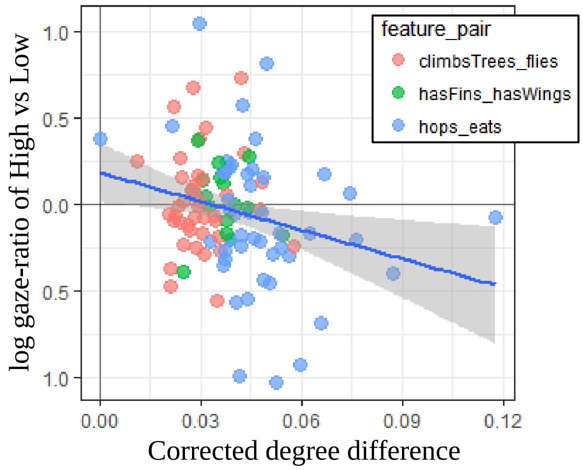

Figure 16. Dot plot of feature Bias log gaze-ratio of High vs Low as a function of the difference in corrected individually calculated degree between high and low items.

The significant negative relation between degree difference and Feature Bias is clearly visible in the dot plot; indicating that for items with a larger difference in individually calculated corrected degree between high and low items participants are more likely to fixate the normative low item across the length of the trial. However, further inspection of the graph uncovers the effect seems to be largely driven by the “hops-eats” pair of features – possibly due to the relatively limited variance in degree difference for the other two pairs of features. Thus, we next included the same coefficients in a linear mixed effects model, with random intercepts for participant and item (for coefficients see Table 9). This model indicated that corrected degree difference is only marginally significantly related to the Feature Bias.

Table 9. Results for the model of log gaze-ratio of fixations to high vs low items in the Feature Bias trial

Variable	Coef.	95% CI	p.val
Intercept	0.010	[-0.374, 0.399]	.95
Visual Bias	0.000	[-0.203, 0.206]	.99
Corrected Deg Dif	**-0.191**	[-0.393, 0.013]	**.06**

Note. Coefficients are β_{std}.

Time window analysis of influence of connectivity on looking to naming exposure videos

We explored how looks to naming exposure videos related to the connectivity characteristics of the features (normative High/Low, corrected individual degree), initial visual biases, order of video presentation, and the three individual differences of interest (age, vocabulary percentile, and attentional focus score). For this analysis, we calculated the proportion of looks over the 13,000 ms naming exposure trial, and constructed a full linear mixed effects model, with random intercepts for participants and feature, in which proportion of looks was predicted as a function of the visual bias (the log gaze-ratio of the proportion of looks to the item over the competitor, over the full 6,000 ms of the trial), the order of video presentation (first or second presentation for a given feature), the two connectivity measures, the three individual differences measures, and the six sets of interactions between the connectivity and individual differences measures. This full model was then fed through a backwards stepwise feature selection algorithm (set to retain random effects). The coefficients for the model output from the backwards selection process are presented in Table 10.[1]

Table 10. Results for the model output from the backwards stepwise feature selection predicting fixation proportions to naming exposure videos.

Variable	Coef.	95% CI	p.val
Intercept	0.008	[-0.203, 0.213]	.94
Order	**-0.237**	[-0.305, -0.17]	**<.001**
Attentional Focus Score	**0.306**	[0.143, 0.469]	**<.001**

Note. Coefficients are β_{std}.

Results indicate a negative effect of order and a positive effect of attentional focus score. In other words, as with feature exposure vides, naming exposure videos are fixated more the first time than the second, and more for participants with higher attentional focus scores.

[1] The same pattern of significant coefficients emerges from backwards selections for models Visual Bias scores calculated over the other two time-windows explored above (i.e., 1,500 and 3,000 ms).

Growth curve analysis of influence of connectivity on looking patterns in the pre-labeling preview period in retention task

To explore whether the qualitative similarity between the timecourses of fixation proportions to normative high and low items in the preview period for novel items here and known items (Chapter 2, this dissertation) is quantitatively supported, we next performed a growth curve analysis (Mirman, Dixon, & Magnuson, 2008; Mirman, 2014) comparing the timecourses of fixation proportions to normative high versus low items using orthogonal polynomial timecodes (i.e., uncorrelated linear, quadratic and cubic components of the time predictor, which are better suited to multiple regression) calculated using eyetrackingR (Dink & Ferguson, 2018). The model consisted of a linear mixed effects model, with random intercepts and slopes for participant and item pair, in which the proportion of fixations to the item was predicted as a function of the normative high vs low status of the item, the first three orthogonal polynomial timecodes, and the interactions between high vs low status and the timecodes. The results of the analysis are presented in Table 11. Here, we simply note that the pattern of significant results perfectly parallels those from the equivalent analysis for known words.

Table 11. Results for growth curve analysis of the timecourse of fixation proportions to normative high vs low items across the preview period.

Variable	Coef.	t.val	p.val
Intercept	**0.352**	55.209	**< .001**
Normative High	**-0.012**	-2.793	**0.005**
OT 1	**0.408**	21.152	**< .001**
OT 2	**-0.359**	-21.154	**< .001**
OT 3	**0.152**	8.494	**< .001**
Normative High: OT 1	**-0.046**	-2.757	**0.006**
Normative High: OT 2	0.029	1.759	0.08
Normative High: OT 3	**0.049**	2.958	**0.003**

Time window analysis of relations between looking patterns during learning opportunities and attention in the post-labeling early and late test period of the retention task

Here we conduct analyses exploring whether performance in the early and late test periods relate to looking patterns in earlier learning trials (i.e., feature and naming exposure trials), and if there are any individual differences in such relations. First, we added mean fixation proportions for feature and naming exposure trials to the previous full linear mixed effects models, along with the six sets of interactions with the three individual differences measures. These full models were then fed through a backwards stepwise feature selection algorithm. While the model output for the early test period was unchanged, that output for the late test period contained additional significant predictors (see Table 12).

Table 12. Results for the model output from the backwards stepwise feature selection of log gaze-ratio of fixations to target vs distractor items in the late test period

Variable	Coef.	95% CI	p.val
Intercept	-0.048	[-0.164, 0.067]	.37
Preview log gaze-ratio	**0.236**	[0.133, 0.34]	**<.001**
Percentile	**0.112**	[0.008, 0.215]	**.04**
MN Naming Exposure Proportion	0.05	[-0.062, 0.163]	.27
Attentional Focus Score	-0.016	[-0.128, 0.096]	.46
MN NE Prop: AF Score	**0.125**	[0.025, 0.226]	**.01**

Note. Coefficients are β_{std}.

The results indicate significant positive main effects of preview period log gaze-ratio to target vs distractor and percentile, and a significant positive interaction between mean proportion of looking time in naming exposure trials and attentional focus scores. We visualize the interaction in Figure 17 by plotting model fit estimates of (scaled) log gaze-ratio of fixations to target vs distractor items in the late test period as a function of mean naming-exposure video fixation proportions and attentional focus scores.

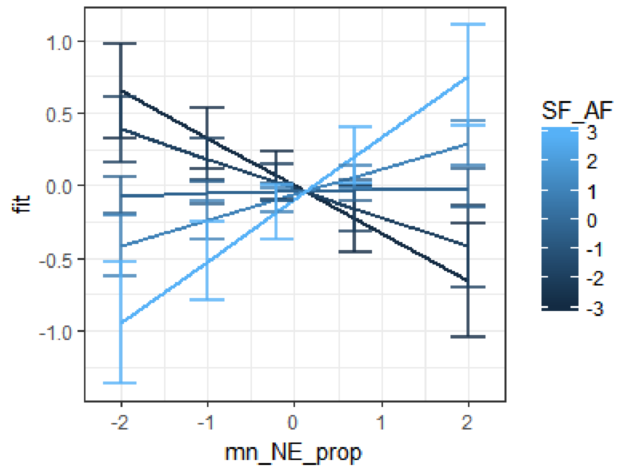

Figure 17. Model estimates (fit) of log gaze-ratio of fixations to target vs distractor items in the late test period as a function of mean naming-exposure video fixation proportions and attentional focus scores.

The positive interaction between mean naming-exposure video fixation proportions and attentional focus scores is clearly visible in Figure 17, which indicates that children with lower attentional scores fixate the target less when they've paid more attention to naming exposure videos, while the reverse is true for children with high attentional focus scores. Taken altogether, the results indicate, in addition to greater looks towards a target in the preview period leading to more looks in the test period, participants who are higher percentile and those who have higher attentional focus scores and paid more attention in naming trials are also more likely to fixate the target in the late test period.

Growth curve analysis of influence of connectivity on looking patterns in the post-labeling test period in the retention task

To explore whether the lack of any significant relations between normative connect and log gaze-ratio of fixations to target vs distractor items was due to the complex non-linear pattern that is visually apparent in the time course plots, we next performed a growth curve analysis comparing the target-distractor log gaze-ratios for normative high and low connectivity conditions paralleling the earlier analysis for the preview period. The results for the analysis are presented in Table 13.

Table 13. Results for growth curve analysis of the timecourse of log gaze-ratios of target versus distractors for normative high and low target conditions across the test period.

Variable	Coef.	t.val	p.val
Intercept	0.018	0.657	0.52
Preview Log Gaze-Ratio	**0.078**	9.868	**< .001**
Normative High	-0.016	-1.021	0.31
OT 1	0.254	1.663	0.11
OT 2	0.024	0.156	0.88
OT 3	-0.178	-1.348	0.18
OT 4	-0.094	-0.825	0.41
Normative High: OT 1	**-0.196**	-2	**0.05**
Normative High: OT 2	**0.226**	2.297	**0.02**
Normative High: OT 3	0.126	1.291	0.20
Normative High: OT 4	**-0.266**	-2.745	**0.006**

Focusing on the coefficients of interest, although the main effect of normative high versus low conditions does not come out significant, the interactions of normative condition with the linear, quadratic and quartic timecodes do significantly contribute to the model. To visualize these effects, in Figure 18 we plot the full model, as well as separate component models of the each of the four interactions.

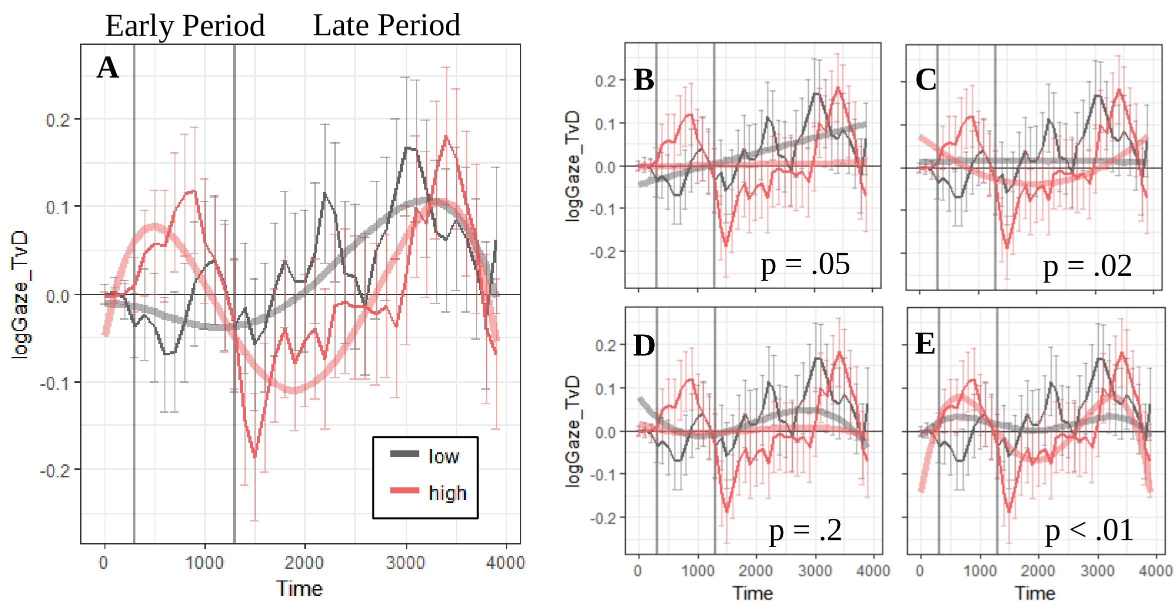

Figure 18. Test period timecourse plots of the log gaze-ratios of target versus distractors for normative high (red) and low (grey) target conditions, overlaid with (A) the full growth curve model, and component models demonstrating the interactions between normative high vs low and the (B) linear, (C) quadratic, (D) cubic, and (E) quartic orthogonal timecode polynomials, calculated over 100 ms time bins (with *SE* bars). Early test period from 300 to 1,300 ms (marked by gray vertical lines), and late test period from 1,300 to 4,000 ms.

The significant negative interaction between normative high vs low target condition and the linear time code characterizes how fixations towards normative high targets are greater at the outset, while fixations towards normative low targets are relatively more robust for the remainder of the trial. The significant positive interaction with the quadratic time code and the negative interaction with the quartic time code both capture how the normative high target condition timecourse has a noticeable pattern of a peak, followed by a valley, ending with a peak; while in contrast the normative low target condition timecourse monotonically rises throughout the trial.

Discussion

The mechanisms underlying individual differences in early vocabulary size remain poorly understood, and are a major focus of ongoing research. One interesting possibility is that individual differences in vocabulary size may relate to differences in vocabulary composition.

This hypothesis is supported by recent findings tying the semantic structure of children's early productive vocabularies to word learning skill on the one hand (e.g., Borovsky et al., 2015) and use of attentional biases on the other (e.g., Colunga & Sims, 2017). Furthermore, these results can be tied together via the Attentional Learning Account (e.g., Colunga & Smith, 2008; Smith & Samuelson, 2006). If children selectively attend to features based on patterns of shared perceptual features between to-be-processed items and known vocabulary, such within object biases could result in between object biases within environments. Furthermore, if children compare the structure of features in novel objects to the correlational structures within the known objects with shared features, they could facilitate the initial creation of representations and aid in initial label learning (Jones et al., 1991). Importantly, individual differences in such processes could snowball over time resulting in large individual differences in both overall vocabulary size and relative familiarity with individual items, as seen in neural network models exploring the role of structure in early lexico-semantic development (Borovsky & Elman, 2006). Building on this previous work and using a recently developed set of semantic features norms for early acquired words (Peters et al, in prep), the primary goal of this project was to explore how the (visual-motion and visual-form and surface) perceptual-semantic connectivity of novel objects influences 24- to 30-month-old toddlers' patterns of attention throughout the process of learning about novel object features and labels, and subsequent object-label learning outcomes. A secondary goal was to determine whether any semantic connectivity effects were related to individual differences in age, word learning skill, or tendency to maintain focused attention.

To probe the influence of connectivity on patterns of attention during opportunities to learn about novel objects' features and labels we ran analyses exploring how the proportion of time attending feature and naming exposure videos vary as functions of (normative and individually calculated) perceptual-semantic connectivity and the three individual differences measures. At the outset, we noted two outcomes of interest: 1) videos demonstrating normative high connectivity videos and/or subsequent labeling videos might draw more looks than low connectivity videos or 2) vice versa. Exploratory analyses of the feature videos revealed positive main effects of attentional focusing score for both feature and naming exposure videos, and a significant negative interaction between attentional focusing and normative connectivity for the feature exposure video. In other words, in addition to paying more attention overall during feature and name learning opportunities, toddlers with high attentional focus scores looked more at videos demonstrating low connectivity features than those demonstrating high connectivity features, while toddlers with low attentional focus scores showed the opposite effect: looking less at videos demonstrating low connectivity features and more to those demonstrating high connectivity features. While this result was somewhat unexpected, one interpretation is that the relative informativity of objects with high versus low semantic connectivity differ for toddlers with high versus low attentional focusing tendencies – with objects with low connectivity features being more informative for toddlers with high attentional focus scores, and objects with high connectivity features being more informative for toddlers with low attentional focus scores. Interestingly, if these looking patterns generalize to real world feature learning opportunities and influence downstream label learning, one hypothesis that emerges from this result is that toddlers with high attentional focus scores may actually have more diverse lexico-semantic networks than toddlers with low attentional focus scores.

To explore how word knowledge connectivity relates to patterns of attentional biases across the experiment we compared looks to high versus low (normative and individually calculated) perceptual-semantic connectivity objects for the three different kinds of opportunities for free viewing (i.e., initial visual bias, post-feature-exposure bias, and post-naming-exposure bias in the preview period of test trials). We noted two, nonexclusive sets of outcomes of interest. First, high or low connectivity features might draw greater attention to objects post feature and/or naming exposure. Second, we could see dynamic changes in biases within a given trial similar to as seen with known words (Chapter 2, this dissertation). Time window analyses of mean looking to normative high versus low items over the post-feature-exposure free viewing trials indicated a marginally significant negative effect of the difference in corrected individually calculated degree between high and low items. In other words, for item pairs with large differences in corrected degree, toddlers looked more towards the normative low connectivity item than for item pairs with small differences. Qualitative and exploratory growth curve analyses of the post-naming-bias time course revealed a pattern similar to the time course for pre-labeling biases seen for known words (this dissertation, Chapter 2) – an early tentative bias towards normative high connect items followed by a slower more robust bias towards normative low connect items. Analyses of mean looking time to normative high versus low items in the preview period of test trials indicated a positive interaction between corrected degree difference and age, indicating the longer-lasting bias towards the normative low degree item is driven by a combination of younger participants paying more attention to normative low degree items for which the difference in individually calculated degree between normative high and low items is relatively larger and older participants paying attention to low degree items in cases where the differences is small. Notably a similar interaction is seen for known words, but between corrected degree and percentile (this dissertation, Chapter 2).

Finally, to investigate the influence of connectivity on novel object-label learning outcomes, and relations with pre-labeling biases, we compared looks to target versus distractor items as a condition of connectivity in the post-labeling test period. Qualitative analyses of the timecourse plots noted evidence for early facilitation for normative high items and later facilitation for normative low connectivity items, paralleling the results for the preview period. Pre-planned time window analyses indicated no robust influence of either normative high/low connectivity or individually calculated corrected degree. However, exploratory growth curve analyses comparing the normative high and low target condition timecourse did detect significant differences in line with the patterns noted in the qualitative analysis, even when considering attentional biases in the preceding preview period. Furthermore, exploratory analyses of the late test period time window also indicated novel-object label learning outcomes related to percentile and an interaction between attentional focusing scores and mean looking times during naming exposure videos, indicating toddlers who are higher percentile and those who have higher attentional focus scores and paid more attention in naming trials are more likely to fixate the target in the late test period.

These results provide three main insights into the fine-grained effects of semantic knowledge on attention to novel objects and cascading effects on subsequent word learning, which in turn provide tentative support for the idea that early object processing is influenced by shared feature-guided selective attention. First and foremost, this work provides evidence that early novel noun-concept processing is influenced by patterns of shared (visual-motion and visual-form and surface) perceptual features between novel objects and known vocabulary from the outset – providing corroborating empirical evidence in line with recent modeling work indicating that such perceptual features matter most in early lexico-semantic development (Peters & Borovsky, 2019). Second, this work provides evidence that, even after just recently learning the featural characteristics and label of a novel object, children demonstrate shared-feature-driven attentional biases similar to those seen for known noun-concept processing (Chapter 2, this dissertation). This point provides support for the idea that the mechanism of shared-feature-guided selective attention can be used to explain the facilitatory effects of category density for both novel object-label learning and known word processing (Borovsky et al., 2015; 2016). Third, the exploratory growth curve analysis of the test period timecourse reveals tentative evidence that novel object-label learning outcomes relate to patterns of perceptual feature connectivity following a similar pattern as for preceding attentional biases (i.e., a fast but fleeting benefit for normative high connect items followed by a slower, but more robust boost for low connect items), but are significant even when considering the immediately preceding attentional biases. This pattern is somewhat similar to as seen for known word processing – indicating either the presence of connectivity effects on label processing separate from preceding attentional biases or that cascading effects of initial biases are nonlinear (e.g., initial biases could inversely relate to early versus later patterns of post-labeling looking).

While this work reveals the above novel insights, there are also a number of limitations that constrain the scope of the findings. First, it is unclear to what degree toddlers are remembering the novel features they are exposed to, and whether certain aspects of the experimental setup or feature videos themselves are facilitating or inhibiting learning. For example, the fact that there is a relatively large gap between feature exposure and test might influence any effects of connectivity. Furthermore, the fact that feature videos are accompanied by audio labels of features could have unintentional consequences. One potential idea would be to combine feature and naming exposure videos, either by incorporating the name into the audio feature descriptions, or to remove audio feature descriptions altogether and simply label the object while visually presenting the feature. Another idea would be to design an experiment specifically to test whether children remember the features, for example by replacing the object-labels in the test trials of current setup with audio feature descriptions. A second limitation of the current study is that the calculations for both normative and individually calculated connectivity assume that the labeled feature from the videos is uniquely driving connectivity. While video pairs were designed to match for any additional features, such matching cannot control for the fact that features that correlate with the demonstrated feature, but are not shown in the video, could be part of the representations that children remember. Indeed, connectionist modeling work (Rogers & McClelland, 2004) suggests such "illusory projections" (Gelman, 1990) result precisely because of the overgeneralization of robustly correlated features to situations where they are not present. One potential way to explore this possibility would be to collect feature norms for the novel objects, a process which could possibly reveal any such hidden correlated features. Finally, the current methods of analyses may not be the most appropriate for capturing the linked patterns of attention throughout the experiment. Future work may benefit from taking an approach specifically designed to explore the cascading influences of effects, such as structural equation modeling.

In conclusion, the work presented in this chapter provides evidence that shared (visual-motion and visual-form and surface) perceptual features between novel objects and known vocabulary items influence attention to novel objects and subsequent object-label learning outcomes throughout the learning process. These results are in close alignment with those from Chapter 2, exploring the influence of connectivity on known word processing, and together provide support for accounts emphasizing the early importance of perceptual features and tentative evidence for feature-guided selective attention as being one of the mechanisms underlying early individual differences in vocabulary size.

CHAPTER 4. CONCLUSION

In the introduction we introduced the research question motivating this project, namely: Does noun-concept processing skill relate to the semantic connectivity of a toddler's vocabulary via attentional biases that depend on relations between the semantic features of to-be-processed noun-concepts and known vocabulary? We noted such a mechanism could partially underly individual differences in early vocabulary size – which robustly relates to language related outcomes in childhood and adolescence (Dale et al., 2003; Lyytinen et al., 2005; Rescorla, 2002;2009), which in turn have cascading effects on later academic achievement, professional attainment, and social and emotional functioning (Bishop & Leonard, 2000; Bronwlie et al., 2004; Conti-Ramsden & Durkin, 2016). We laid the groundwork for this hypothesis by first covering work exploring relations between the structure of toddlers' word meaning knowledge and their word processing skill (e.g., Borovsky et al., 2015; 2016), before then introducing work relating vocabulary composition to the use of attentional word learning biases (e.g., Colunga & Sims, 2017). We then tied these bodies of work together via an application of the Attentional Learning Account (Smith & Samuelson, 2006) to describe shared-feature-guided selective attention as an explanation of facilitatory effects of high-category density on novel and known word processing. A key point of the extension was that, rather than a broadly applicable mechanism along the lines of the shape bias, such shared-feature-guided processes should be more idiosyncratic to clusters of known objects that form categories. In other words, to explore the issue we needed to move beyond simple descriptions of categories or features, and work with more granular featural descriptions of objects to explore patterns of shared features across the full breadth of individual toddlers' vocabularies. Therefore, we introduced the methodology of modeling toddlers' vocabularies using noun-feature networks – an approach that builds on recent advances in the application of graph-theoretic network modeling techniques to lexico-semantic representations (e.g., Hills et al., 2009a, 2009b; Steyvers, & Tenenbaum, 2005). Furthermore, building on recent modeling work showing perceptual features matter most in early lexico-semantic development (Peters & Borovsky, 2019), we focused our attention on (visual-motion and visual-form and surface) perceptual-semantic features of early learned noun-concepts according to a set of feature production norms for early learned nouns (Peters, McRae, & Borovsky, in prep).

Altogether, the work presented in the introduction demonstrated the utility of noun-feature networks for exploring developmental questions, and clearly laid the groundwork for a way forward. Namely, empirical research was necessary to explore whether shared perceptual features truly influence early word processing, and if so, whether shared-feature-driven shifts in selective attention are a possible mechanism. Thus, the bulk of this dissertation was dedicated to presenting two eye-tracking studies that empirically explored whether early word processing and learning is indeed facilitated by shared perceptual features with other concepts. In chapter 2, we presented work investigating how patterns of perceptual connectivity in the productive vocabularies of toddlers influence their initial attentional biases to *known* objects and subsequent label processing. In chapter 3, we presented the results of a word-learning task likewise designed to investigate how patterns of shared perceptual features influence attentional biases to *novel* objects and subsequent learning outcomes of object-label pairings.

In this concluding chapter we tie together the results described in the preceding chapters into a coherent whole. We begin by summarizing the key results, then discuss whether they provide support for the hypothesized mechanism of shared-feature-guided selective attention, before finally turning towards the future.

Key Findings

The key findings that we would like to emphasize are:

- Patterns of shared (visual-motion and visual-form and surface) perceptual features relate to differences in early noun-concept processing at the individual level.
- Such influences are at play from the outset of novel noun-concept learning.
- Connectivity driven attentional biases to both recently learned and well-known objects follow a similar timecourse and show similar patterns of individual differences.
- Initial, pre-labeling attentional biases to objects relate to subsequent label processing, but do not linearly explain effects of connectivity.

We next consider each of these points in greater depth in turn.

Shared perceptual features matter for early noun-concept processing at the individual level

As previously noted, recent modeling work has shown that, at the *group* level, patterns of shared perceptual features between early-learned noun-concepts drive lexico-semantic development in 16- to 30-month-olds (Peters & Borovsky, 2019). However, the question remained: Do patterns of perceptual similarity within *individual* toddlers' noun-concept vocabularies likewise influence their online lexico-semantic processing? The results of the studies described in Chapters 2 and 3 revealed that the answer is: Yes. For both well-known and recently learned novel noun-concepts, analyses revealed influences of perceptual connectivity on both attentional biases and label processing. These results also provide insight into another question raised in the aforementioned modeling work. Namely, as discussed in the introduction, while the modeling work revealed a relation between number of shared perceptual features and the normative order in which toddlers learn words, analyses also supported the possibility that the effect was actually the result of a robust lower-level effect (i.e., number of features) propagating up to higher levels. In other words, it was possible that patterns of similarity structure did not actually play a driving role in determining the order of early word learning and lexico-semantic network structure, but instead arose out of the ease of individual word learning. If this was true, it would arguably call into question the entire mechanism of shared-feature-guided shifts in selective attention as a possible underlying force in early word learning. However, the results from this project indicate that, even when controlling for the number of features, patterns of shared perceptual features still influence early noun-concept processing.

Shared perceptual features influence noun-concept learning from the outset

In chapter 3, exploratory analyses revealed influences of perceptual connectivity on patterns of attention to videos demonstrating features of novel objects. Specifically, an interaction effect (with attentional focus score) provided evidence that toddlers pay more attention to videos demonstrating rarer, low connectivity features. These videos were participants' first exposure to information about the objects beyond their visual appearance, which they first saw in the immediately preceding trial. In other words, this result provides evidence that perceptual connectivity influences noun-concept learning from nearly the earliest point possible. This effect was in the opposite direction of what would be expected if the recognition of shared features draws attention, instead hinting that relatively novel featural characteristics may draw greater attention in the earliest stages of learning about a novel object.

Connectivity driven attentional biases to newly learned and well-known objects are similar

Qualitative analyses of the timecourse plots and exploratory growth curve analyses of the pre-labeling preview period for both known and recently learned noun-concepts revealed the same pattern: participants briefly pay relatively greater attention to normative high connectivity items first, before then shifting their attention to the normative low connectivity items for the remainder of the trial. In other words, attention to high connectivity items is fast and first, but brief, while attention to low connectivity items is slow and second, but robust. Furthermore, analyses revealed parallel interactions between connectivity and individual differences measures for known and novel items. For known words, exploratory time window analyses revealed an interaction between corrected degree difference and percentile, indicating the slow and second bias towards low connectivity items was driven by high percentile participants paying more attention to low degree items for which the difference between normative low and high connect items is relatively small. For novel words, planned time window analyses revealed a similar interaction, but with age. These similar patterns of attentional biases provide tentative evidence that well-developed and recently formed noun-concept representations are visually processed in qualitatively similar ways, in line with what might be predicted based on recent work placing known and novel noun-concept representations on a single continuum (McMurray, Horst, & Samuelson, 2012).

Initial biases relate to label processing, but do not linearly explain effects of connectivity

For both known and novel noun-concepts, patterns of attention when processing labels show parallels with preceding pre-labeling attentional biases. For known words, qualitative analysis of the timecourse and planned time window analyses of the test period indicated that label processing is facilitated for normative high connectivity items, reflecting the pattern of fast and first attention to the same items in the preview period. For novel words, qualitative analysis of the time course and exploratory growth curve analysis revealed a fast but fleeting benefit for normative high connect items followed by a slower, but more robust boost for low connect items, paralleling the results for the preview period. While pre-labeling attentional biases related to subsequent label processing, the relations between connectivity and both known and novel label processing were significant even when taking into account pre-labeling biases – indicating either the presence of additional effects of connectivity on auditory label processing or that cascading effects of initial biases on subsequent processing are nonlinear and dynamic.

Is there Support for Shared-Feature-Guided Selective Attention?

The mechanism of shared-feature-guided selective attention came into focus as a possible mechanism underlying early differences in toddler vocabulary growth in the introduction, when we tried to use the Attentional Learning Account (e.g., Smith, 2000; Smith & Colunga, 2008; Smith & Samuelson, 2006) to explain facilitatory effects of high category density in individual toddlers' productive vocabularies on both known object-label processing and novel object-label learning outcomes (Borovsky et al., 2015; 2016). We hypothesized that, if such a mechanism were to be at play, we would expect patterns of shared features in individual toddlers' vocabularies to relate to their online processing of noun-concepts. Altogether, the key findings described above do offer tentative support for the idea that early noun-concept processing is influenced by feature-guided selective attention related to patterns of shared (visual-motion and visual-form and surface) perceptual features between to-be-processed items and known vocabulary. The most straightforward support comes from the finding that patterns of shared features relate to early noun-concept processing even when accounting for word level characteristics that have been shown to influence such processing, including number of features and frequency in child directed speech. Additional support comes from the parallel results for the pre-labeling preview periods and post-labeling test periods for both known and recently learned items. Through the lens of this mechanism, the initial bias towards both high connectivity items in the preview period can be explained as resulting from the recognition of shared perceptual features drawing attention. Moving to the test period, the facilitatory effect of high connectivity items for known words could be explained by an additional processing boost that results from a history of relatively greater attention to the high connectivity items, while the early boost for recently learned words could be explained as facilitated object-label learning resulting from more robust object representations.

While the above results can be used to support the role of shared-feature-guided selective attention as a key mechanism in early noun-concept processing, certain aspects of the results raise complications for such an interpretation and necessitate elaboration. First, while the fact that shared features influence noun-concept learning from the outset does provide support for some kind of mechanism involving the recognition of shared features, the fact low connectivity feature videos drew greater attention than high connectivity videos at the very least indicates initial novel object salience could be inversely related to number of shared features. However, such an initial attentional boost for distinctive items does not necessarily argue against processing benefits for items with more shared features once they have been attended – indeed, these two mechanisms could be complementary. For example, if highly connected items benefit from a shared-feature-guided selective attentional boost to representation construction from the outset, the salience of such items when presented in isolation could be relatively short-lived compared to more distinctive items, given the relative lack of new information. A second complication comes from the pattern of biases in the preview period, wherein attention to high connectivity items was fast and first, but fleeting, while attention to low connectivity items was slow and second, but robust. However, this too could be explained as resulting from the relative ease of processing high connectivity items, which would allow for greater attentional resources dedicated to more informative, but difficult to process low connectivity items.

In this dissertation we've focused on the Attentional Learning Account because of its combined focus on the learning of both the featural aspects of concepts and their labels. However, we'd like to briefly make note of Rakison and colleagues' proposal for *constrained attentional associate learning* (CAAL; Rakison, 2005, 2006; Rakison, Lupyan, Oakes, & Walker-Andrews, 2008) – a similar theoretical framework that focuses on earlier conceptual development, and was specifically designed to explain how infants develop representations capable of distinguishing animate from inanimate noun-concepts. As with the ALA, the CAAL framework is grounded in the idea that the primary mechanisms driving early (lexico-)semantic development are domain-general associative learning in combination with attentional constraints on online processing and resultant learning that emerge due to past experience with statistical regularities across features in the input. Furthermore, the hypothesized attentional constraints are likewise guided by features that are shared between known concepts and the current input. However, the CAAL framework has additional stipulations missing from the ALA. First, given the focus on the animate-inanimate distinction, the CAAL framework emphasizes the differences between 'static' and 'dynamic' cues. Where static cues share much in common with the visual-form and surface perceptual features described in this dissertation, while dynamic cues are nearly corollary with our visual-motion features. However, while in this dissertation we have collapsed these two types of features into simply being the relevant perceptual features, Rakison and colleagues have emphasized the differences between static and dynamic cues. They have noted that static cues are immediately apparent in even a single frame of input, whereas dynamic cues are "available only intermittently in the perceptual input, which means that infants will learn about them more slowly than about static ones… such cues are inherently more difficult to process than static ones precisely because their dynamic nature makes them more complex" (Rakison et al., 2008, p. 4). Second, the CAAL framework highlights how maturational differences in information-processing abilities, such as short- and long-term memory and processing speed, relate to the development of object concepts in combination with experience and the aforementioned differences between static and dynamic object features. Thus, while the results presented in this dissertation leave the main proposal (that shared-feature-guided selective attention is a key mechanism underlying individual differences in early lexico-semantic development) underdetermined, the CAAL framework provides additional support and highlights possible future directions – which we discuss in the following section.

Looking Forward

Before discussing possible future directions, let us first reiterate the cascade of processes of interest in more generalizable language. To set the stage, we have individual toddlers, with noun-concept knowledge resulting from unique experiential trajectories. This knowledge includes the implicit recognition of patterns of correlations across various object, linguistic, and contextual features, and serves as a constantly evolving individualized context brought to any given processing opportunity. At the broadest level, we are then interested in whether there are any patterns of differences in such individualized contexts that relate to individual differences in language skills. Here, we more narrowly focused on how the semantic featural structure of such individualized contexts might influence the lexico-semantic processing of individual noun-concepts. Limiting our focus as such in turn allowed us to extend our attention along the developmental trajectories of these individual noun-concepts – from first exposure, through initial object feature and label learning, to robust noun-concept representation processing. In this section, we work our way backwards through this sequence, reconsidering what we know in light of the results of the current study, note apparent gaps in our understanding, and discuss possible future directions.

Beginning with known word processing, the results of the study in Chapter 2 indicate that patterns of shared perceptual features influence both visual attention to objects and auditory object-label processing. Furthermore, the pattern of results generally aligns with what might be predicted based on the combination of Borovsky and colleagues (2016) finding of a facilitatory effect for high density category knowledge on known word processing and Peters and Borovsky's (2019) finding that perceptual features matter most for early lexico-semantic development. However, due to a number a limitations, much remains unclear. First, in the current study, the fact that item categories were limited to vehicles and animals and that features were limited to (visual-motion and visual-form and surface) perceptual features, limits the degree to which we can generalize these results to other categories on the one hand, and doesn't clarify whether we might see similar patterns when considering other types of features on the other. Indeed, researchers have noted that various types of categories, such as living things versus artifacts, differ in the structure of their featural makeup – with the representations of living things tending to consist of clusters of correlated features, while featural representations for artifacts (i.e., things made for humans to be used by humans) generally center on the intended function of the creator (e.g., Gelman, 1988; McRae et al., 1997). Considering the known noun-concepts in the current study included both living things and artifacts (animals and vehicles, respectively), collapsing across categories may have actually introduced noise and weakened any effects. Thus, future work may benefit from not only including a wider variety of categories and feature types, but also by considering how varying along these two dimensions may interact. Finally, the significant effect of presentation order on the relation between preview period attentional biases and test period label processing reveals that effects of connectivity on attentional biases and processing are influenced either by the recency with which noun-concepts have last been processed or the temporal context of other noun-concepts. One potential way to explore this issue would be to employ computational modeling – with one useful example being Mirman and Magnuson's (2008, 2009) use of Cree and colleagues' (Cree, McRae, & McNorgan, 1999) attractor dynamical model to explore the role of semantic feature based similarity structure in known word lexico-semantic processing.

Moving on to the processing of recently learned noun-concepts, the findings of Chapter 3 likewise indicated that patterns of shared perceptual features influence both visual attention to objects and auditory object-label processing. As noted above, the close similarity between findings for well-known and recently learned noun-concepts is in line with what might be predicted based on recent work placing known and novel noun-concept representations on a single continuum (McMurray, Horst, & Samuelson, 2012). However, as with known words, the generalizability of the findings is limited due to the facts that all items were assigned features that generally associate with animals and that features were limited to (visual-form and surface and visual-motion) perceptual features. Furthermore, while the alignment in findings supports the inference that participants were treating the recently learned noun-concepts as "real" noun-concepts, future work employing this paradigm may benefit from more direct exploration of the characteristics of the representations participants are forming. For example, it may be worth designing an experiment to specifically test whether children remember the features, perhaps by replacing the object-labels in the test trials of current setup with audio feature descriptions.

The arguably most novel aspect of the research presented in this dissertation was the feature exposure portion of the novel word study in Chapter 3. The findings provided tentative evidence that patterns of shared perceptual features influence noun-concept processing from nearly the earliest point possible. However, when considering how this task compares to equivalent real-world feature learning opportunities, a number of limitations become apparent. The most obvious limitation is that toddlers generally don't learn about individual featural characteristics of objects through ten-second-long cartoons – instead they have more extended, multi-modal exposure to a variety of features, during which they typically have some degree of agency over their actions. For example, imagine a child that sees a goat for the first time in a petting zoo. After seeing, hearing, and maybe even smelling the goat from a distance, she may have the opportunity to walk up to it and touch it, and see how it responds to her and other stimuli in the environment. Such learning opportunities contain complex arrays of semantic, linguistic, and contextual features, with structures that relate to an individual's mental context far more intricately than the videos we used in the current paradigm, which in turn possibly results in representations that are far more robust. In other words, while the current task setup allows us to control the stimuli for variety of variables, it also may be weakening the very effects we are trying to explore. While incorporating real-world feature exposure opportunities into experiments should be a long-term goal, there are a couple of intermediary steps that are more accessible. For example, rather than a single feature, participants could be exposed to an array of features. Or, rather than being exposed in only one session, participants could be exposed to features across a series of sessions. It may also be possible to include some degree of agency by using a fixation contingent looking paradigm to allow participants to choose what objects and features they are exposed to. Finally, we could begin to take steps towards realistic feature exposure, by presenting participants with real objects that they can see, touch and play with. While this last approach is relatively common in the categorization literature, the key would be to design objects that can be variably assigned different features from the feature production norms, or select of set of actual novel objects that can be balanced for other psycholinguistic factors that may influence learning.

Next, it is worth taking a moment to consider how our findings may have been influenced by our methods for characterizing lexico-semantic knowledge and structure. We chose to use graph-theoretic noun-feature networks because the methodology allowed us to move beyond simple descriptions of categories or feature dimensions, and work with more granular patterns of shared features. This method clearly aligns with the research question driving this project, but some of the implementation decisions – in particular regarding network construction and choice of metrics – may be worth reconsidering. First, when building our networks, following prior precedent (e.g., Engelthaler & Hills, 2017; Hills et al., 2009a), links were treated as unweighted, undirected edges. This means that any variety of patterns of shared features between any two concepts were collapsed into a single, uniform link. For example, a link between a pair of concepts that share a single feature was treated as equivalent to a link between a pair of concepts that shared a number of highly correlated features. However, McRae and colleagues (1997) showed that such differences in pairwise similarity structure matters: they found that shared individual features relate to priming for artifacts but not living things, while sets of shared correlated features relate to priming for living things but not artifacts. Furthermore, it is also worth noting that all types of features were treated equally. However, as noted in the previous section, there is reason to believe that some types of features (e.g., static versus dynamic) may be harder or easier to process and learn (Rakison et al., 2008). Thus, future work may benefit from exploring how such differences in pairwise similarity structure may influence how connectivity relates to patterns of attention and processing. Second, we chose to focus exclusively on degree in this project for the sake of simplicity and coherence, however the theory motivating our research questions hints that it may be worth exploring other metrics in future work. For example, based on the explanation of shared-feature-guided selective attention, we might also expect to see differences in processing for concepts depending on the correlational structure of the features making up both the concepts themselves and the neighbors of concepts. Thus, this work should be taken as a relatively simple, foundation building step towards future explorations that make fuller use of the capabilities of the graph-theoretic noun-feature network modeling methodology.

Finally, in the current project, we narrowed our focus to exploring how the semantic featural structure of an individual’s word knowledge might influence the lexico-semantic processing of individual noun-concepts. However, we are interested in placing this work in the broader framework asking: What patterns of differences in experience-dependent processing contexts relate to individual differences in language skills? In the current work, we moved somewhat in this broader direction by including explorations of three individual difference measures: age, vocabulary size percentile, and temperamental tendency to maintain focused attention – which served to capture individual differences in maturational characteristics and breadth of experience, language learning skills, and attentional skills, respectively. These individual differences were only of secondary interest in this work – limiting the rigor and thoroughness with which we were able to address them, but their inclusion served to set the groundwork for more focused future explorations. Indeed, one motivation for this project was to prepare for studies exploring connectivity effects in the processing of children at the lower end of the vocabulary size spectrum, with the goal of eventually developing targeted interventions based on individual children's vocabulary structures. However, before working on any interventions, it seems appropriate to replicate the current study with slightly older children, including a group who have been diagnosed with developmental language disorder, along with age and language matched peers. Future work may also benefit from expanding beyond the current word learning paradigm, perhaps to include greater ambiguity – as found in real-world word learning opportunities. One idea would be to merge the current feature exposure paradigm with the cross-situational word-learning task (e.g., Smith & Yu. 2008) – which was designed to explore how infants learn object-label pairs from co-occurrences of object and labels in noisy, ambiguous contexts. Such a combination would bridge the current work with the broader literature on experience dependent cuing of attention – of which shared-feature-guided selective attention is just one example. Such cued attention has been noted as a central constraint that determines what aspects of early word learning experiences make contact with an infant's learning system (Smith, Suanda, & Yu, 2014), with the potential to either facilitate or obstruct early lexico-semantic learning depending on complex interactions of predictive cues across experience (e.g., Smith, Colunga, & Yoshida, 2010). Thus, merging these two paradigm seems like one particularly promising path towards shedding light on whether differences in individuals' lexico-semantic knowledge can variously guide or hinder young language learners on their journeys through their own unique experiential landscapes.

Conclusion

The most robust finding of this project is that patterns of shared (visual-motion and visual-form and surface) perceptual features related to both initial attention to object images and subsequent processing of auditory labels of both known and recently learned noun-concepts. This finding in turn provides support for accounts emphasizing the early importance of perceptual features and tentative evidence for feature-guided selective attention as being one of the mechanisms underlying early individual differences in vocabulary size. While there are numerous limitations with this work, it provides groundwork for more focused explorations with at-risk children at the lower end of the vocabulary spectrum, which in turn may eventually pave the way for targeted interventions based on individual children's vocabulary structures.

REFERENCES

Arai, M., Van Gompel, R. P., & Scheepers, C. (2007). Priming ditransitive structures in comprehension. *Cognitive psychology*, *54*(3), 218-250.

Arias-Trejo, N., & Plunkett, K. (2013). What's in a link: Associative and taxonomic priming effects in the infant lexicon. *Cognition*, 128(2), 214-227.

Barabási, A. L. (2016). *Network science*. Cambridge University Press.

Barsalou, L. W. (2003). *Abstraction in perceptual symbol systems*. Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences, 358, 1177–1187.

Bates, E., & MacWhinney, B. (1987). Competition, variation, and language learning. In B. MacWhinney (Ed.), *Mechanisms of language acquisition*, Hillsdale, NJ: L Erlbaum Associates.

Bates, D., Maechler, M., Bolker, B., Walker, S., Christensen, R. H. B., Singmann, H., ... & Rcpp, L. (2015). Package 'lme4'. *Convergence*, *12*(1).

Beckage, N., Aguilar, A., & Colunga, E. (2015). Modeling lexical acquisition through networks. In *Proceedings of CogSci*.

Beckage, N., Smith, L., & Hills, T. (2011). Small worlds and semantic network growth in typical and late talkers. *PloS one*, *6*(5), e19348.

Bishop, D. V. M., & Leonard, L. B. (Eds.). (2000). *Speech and language impairments in children: Causes, characteristics, intervention and outcome*. New York, NY: Psychology Press.

Borovsky, A., Ellis, E. M., Evans, J. L., & Elman, J. L. (2015). Lexical leverage: category knowledge boosts real-time novel word recognition in 2-year-olds. *Developmental science*.

Borovsky, A., Ellis, E. M., Evans, J. L., & Elman, J. L. (2016). Semantic structure in vocabulary knowledge interacts with lexical and sentence processing in infancy. *Child Development*, *87*(6), 1893-1908.

Borovsky, A., & Elman, J. (2006). Language input and semantic categories: A relation between cognition and early word learning. *Journal of child language*, *33*(4), 759-790.

Braginsky, M., Yurovsky, D., Marchman, V. A., & Frank, M. C. (2016, August). *From uh-oh to tomorrow: Predicting age of acquisition for early words across languages.* Paper presented at the 38th Annual Meeting of the Cognitive Science Society, Philadelphia, PA.

Brownlie, E. B., Beitchman, J., Escobar, M., Young, A., Atkinson, L., Johnson, C. (2004). Early Language Impairment and Young Adult Delinquent and Aggressive Behavior. *Journal of Abnormal Child Psychology*, 32(4), 453-467.

Cohen LB, Cashon Cara H. Infant cognition. In: Kuhn D, Siegler RS, Damon W, Lerner RM, editors. Handbook of Child Psychology: Vol 2, Cognition, Perception, & Language. 6. Hoboken, NJ: Wiley; 2006. pp. 214–51.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological review*, *82*(6), 407.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior*, *8*(2), 240-247.

Collisson, B. A., Grela, B., Spaulding, T., Rueckl, J. G., & Magnuson, J. S. (2015). Individual differences in the shape bias in preschool children with specific language impairment and typical language development: Theoretical and clinical implications. *Developmental science, 18*(3), 373-388.

Colunga, E., & Sims, C. (2011). Early talkers and late talkers know nouns that license different word learning biases. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 33, No. 33).

Colunga, E., & Sims, C. E. (2017). Not Only Size Matters: Early-Talker and Late-Talker Vocabularies Support Different Word-Learning Biases in Babies and Networks. *Cognitive science*, *41*, 73-95.

Colunga, E., & Smith, L. B. (2008). Knowledge embedded in process: the self-organization of skilled noun learning. *Developmental Science, 11*(2), 195-203.

Conti-Ramsden, G., & Durkin, K. (2016). Language impairment and adolescent outcomes. In K. Durkin, & Schaffer, H. R., (Ed.), *The Wiley Handbook of Developmental Psychology in Practice: Implementation and Impact*. West Sussex, UK: Wiley-Blackwell.

Cree, G. S., McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science*, *23*(3), 371-414.

Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, *132*(2), 163.

Dale, P. S., Price, T. S., Bishop, D. V. M., & Plomin, R. (2003). Outcomes of early language delay: I. Predicting persistent and transient language difficulties at 3 and 4 years. *Journal of Speech, Language & Hearing Research*, 46(3), 544-560.

De Deyne, S., Peirsman, Y., & Storms, G. (2009, January). Sources of semantic similarity. In *Proceedings of the 31th Annual Conference of the Cognitive Science Society* (pp. 1834-1839).

Dink, J. W., & Ferguson, B. (2018). eyetrackingR: An R library for eye-tracking data analysis. *Available at www. eyetracking-r. com. Accessed October*, 2018.

Dixon Jr, W. E., & Shore, C. (1997). Temperamental predictors of linguistic style during multiword acquisition. *Infant Behavior and Development*, *20*(1), 99-103.

Dixon Jr, W. E., & Smith, P. H. (2000). Links between early temperament and language acquisition. *Merrill-Palmer Quarterly (1982-)*, 417-440.

Engelthaler, T., & Hills, T. T. (2016). Feature biases in early word learning: network distinctiveness predicts age of acquisition. *Cognitive science*.

Fenson L, Marchman VA, Thal D, Dale PS, Reznick JS, Bates E. (2006). MacArthur-Bates Communicative Development Inventories: User's Guide and Technical Manual. 2nd Ed. Brookes Publishing; Baltimore, MD.

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 et al., 2008

Fernald, A., Zangl, R., Portillo, A. L., & Marchman, V. A. (2008). Looking while listening: Using eye movements to monitor spoken language comprehension by infants and young children. In I. A. Sekerina, Fernandez, E.M., & Clahsen, H. (Ed.), *Developmental Psycholinguistics: On-line methods in children's language processing*. Philadelphia: John Benjamins Publishing Company.

Fox, J. (2003). Effect displays in R for generalised linear models. *Journal of statistical software*, *8*(15), 1-27.

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.

Gelman, R. (1990). First principles organize attention to and learning about relevant data: Number and the animate-inanimate distinction as examples. *Cognitive science*, *14*(1), 79-106.

Gelman, S. A. (1988). The development of induction within natural kind and artifact categories. *Cognitive psychology*, *20*(1), 65-95.

Gershkoff-Stowe, L., & Smith, L. B. (2004). Shape and the first hundred nouns. *Child development*, *75*(4), 1098-1114.

Gilhooly, K. J., & Logie, R. H. (1980). Age-of-acquisition, imagery, concreteness, familiarity, and ambiguity measures for 1,944 words. *Behavior Research Methods & Instrumentation*, *12*(4), 395-427.

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.

Griffiths, T. L., Steyvers, M., & Firl, A. (2007). Google and the mind: Predicting fluency with PageRank. *Psychological Science*, *18*(12), 1069-1076.

Grondin, R., Lupker, S. J., & McRae, K. (2009). Shared features dominate semantic richness effects for concrete concepts. *Journal of Memory and Language*, *60*(1), 1-19.

Hampton, J. A. (1997). Conceptual combination: Conjunction and negation of natural concepts. *Memory & cognition*, *25*(6), 888-909.

Hills, T. (2013). The company that words keep: comparing the statistical structure of child- versus adult-directed language. *Journal of child language*, *40*(03), 586-604.

Hills, T. T., Maouene, M., Maouene, J., Sheya, A., & Smith, L. (2009a). Longitudinal Analysis of Early Semantic Networks Preferential Attachment or Preferential Acquisition? *Psychological Science*, 20(6), 729-739.

Hills, T. T., Maouene, M., Maouene, J., Sheya, A., & Smith, L. (2009b). Categorical structure among shared features in networks of early-learned nouns. *Cognition*, *112*(3), 381-396.

Hills, T. T., Maouene, J., Riordan, B., & Smith, L. B. (2010). The associative structure of language: Contextual diversity in early word learning. *Journal of memory and language*, *63*(3), 259-273.

Horst, J. S., & Hout, M. C. (2016). The Novel Object and Unusual Name (NOUN) Database: A collection of novel images for use in experimental research. *Behavior research methods*, *48*(4), 1393-1409.

Imai, M., & Gentner, D. (1997). A cross-linguistic study of early word meaning: Universal ontology and linguistic influence. *Cognition*, *62*(2), 169-200.

Imai, M., Gentner, D., & Uchida, N. (1994). Children's theories of word meaning: The role of shape similarity in early acquisition. *Cognitive development*, *9*(1), 45-75.

Inhelder, B., & Piaget, J. (1964). *The growth of logic in the child*. London UK: Routledge & Kegan Paul.

Jones, S. S. (2003). Late talkers show no shape bias in a novel name extension task. *Developmental Science*, *6*(5), 477-483.

Jones, S. S., & Smith, L. B. (1998). How children name objects with shoes. *Cognitive Development*, *13*(3), 323-334.

Jones, S. S., Smith, L. B., & Landau, B. (1991). Object properties and knowledge in early lexical learning. *Child development*, *62*(3), 499-516.

Kannass, K. N., & Oakes, L. M. (2008). The development of attention and its relations to language in infancy and toddlerhood. *Journal of Cognition and Development*, *9*(2), 222-246.

Knoeferle, P., & Kreysa, H. (2012). Can speaker gaze modulate syntactic structuring and thematic role assignment during spoken sentence comprehension?. *Frontiers in Psychology*, *3*, 538.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13).

Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive development*, *3*(3), 299-321.

Landau, B., Smith, L. B., & Jones, S. (1992). Syntactic context and the shape bias in children's and adults' lexical learning. *Journal of Memory and Language*, *31*(6), 807-825.

Landau, B., Smith, L., & Jones, S. (1998). Object shape, object function, and object name. *Journal of memory and language*, *38*(1), 1-27.

Lyytinen, P., Eklund, K., & Lyytinen, H. (2005). Language development and literacy skills in late-talking toddlers with and without familial risk for dyslexia. *Annals of dyslexia*, *55*(2), 166-192.

MacWhinney, B. (1987). The competition model. In B. MacWhinney (Ed.), *Mechanisms of language acquisition*, Hillsdale, NJ: L Erlbaum Associates.

MacRoy-Higgins, M., & Montemarano, E. A. (2016). Attention and word learning in toddlers who are late talkers. *Journal of child language*, *43*(5), 1020-1037.

Madole, K. L., & Oakes, L. M. (1999). Making sense of infant categorization: Stable processes and changing representations. *Developmental Review*, *19*(2), 263-296.

Maki, W. S., & Buchanan, E. (2008). Latent structure in measures of associative, semantic, and thematic knowledge. *Psychonomic Bulletin & Review*, *15*(3), 598-603.

Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive psychology*, *20*(2), 121-157.

McMurray, B., Horst, J. S., & Samuelson, L. K. (2012). Word learning emerges from the interaction of online referent selection and slow associative learning. *Psychological review*, *119*(4), 831.

McRae, K., Cree, G. S., Westmacott, R., & De Sa, V. R. (1999). Further evidence for feature correlations in semantic memory. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *53*(4), 360.

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37(4), 547-559.

McRae, K., De Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*(2), 99.

Medin, D. L. (1989). Concepts and conceptual structure. *American psychologist*, *44*(12), 1469.

Mirman, D. (2017). *Growth curve analysis and visualization using R*. Chapman and Hall/CRC.

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.

Mirman, D., & Magnuson, J. S. (2008). Attractor dynamics and semantic neighborhood density: processing is slowed by near neighbors and speeded by distant neighbors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(1), 65.

Mirman, D., & Magnuson, J. S. (2009). Dynamics of activation of semantically similar concepts during spoken word recognition. *Memory & cognition*, *37*(7), 1026-1039.

Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of child language*, *17*(2), 357-374.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of experimental psychology: general*, *106*(3), 226.

Neely, J. H. (1991) Semantic priming effects in visual word recognition: A selective review of current findings and theories. In: *Basic processes in reading: Visual word recognition*, ed. D. Besner & G. W. Humphreys. Erlbaum.

Newman, R., Samuelson, L., & Gupta, P. (2008). Learning novel neighbors: Distributed mappings help children and connectionist models. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 30, No. 30).

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., & Samuelson, L. K. (2011). The shape of the vocabulary predicts the shape of the bias. *Frontiers in Psychology*, *2*, 345.

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.

Peters, R., & Borovsky, A. (2019). Modeling early lexico-semantic network development: Perceptual features matter most. *Journal of Experimental Psychology: General*, *148*(4), 763.

Peters, R., McRae, K., & Borovsky, A. (in prep). An extension of the McRae et al. (2005) semantic feature production norms to cover early-learned concepts.

Petersen, S. E., & Posner, M. I. (2012). The attention system of the human brain: 20 years after. *Annual review of neuroscience*, *35*, 73-89.

Pexman, P. M., Hargreaves, I. S., Siakaluk, P. D., Bodner, G. E., & Pope, J. (2008). There are many ways to be rich: Effects of three measures of semantic richness on visual word recognition. *Psychonomic Bulletin & Review*, *15*(1), 161-167.

Pexman, P. M., Lupker, S. J., & Hino, Y. (2002). The impact of feedback semantics in visual word recognition: Number-of-features effects in lexical decision and naming tasks. *Psychonomic Bulletin & Review*, *9*(3), 542-549.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual review of neuroscience*, *13*(1), 25-42.

Pruden, S. M., Hirsh-Pasek, K., Golinkoff, R. M., & Hennon, E. A. (2006). The birth of words: Ten-month-olds learn words through perceptual salience. *Child development*, *77*(2), 266-280.

Putnam, S. P., Gartstein, M. A., & Rothbart, M. K. (2006). Measurement of fine-grained aspects of toddler temperament: The Early Childhood Behavior Questionnaire. *Infant behavior and development*, *29*(3), 386-401.

Quillian, M. R. (1967). Word concepts: A theory and simulation of some basic semantic capabilities. *Behavioral science*, *12*(5), 410-430.

Quinn, P. C., & Eimas, P. D. (1997). A reexamination of the perceptual-to-conceptual shift in mental representations. *Review of General Psychology*, *1*(3), 271-287.

R Core Team (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org/.

Rabovsky, M., Sommer, W., & Abdel Rahman, R. (2012). The time course of semantic richness effects in visual word recognition. *Frontiers in Human Neuroscience, 6*, 11.

Rakison, D. H. (2005). A secret agent? How infants learn about the identity of objects in a causal scene. *Journal of Experimental Child Psychology*, *91*(4), 271-296.

Rakison, D. H. (2006). Make the first move: how infants learn about self-propelled objects. *Developmental Psychology*, *42*(5), 900.

Rakison, D. H., Lupyan, G., Oakes, L. M., & Walker-Andrews, A. S. (2008). Developing object concepts in infancy: An associative learning perspective. *Monographs of the Society for Research in Child Development*, i-127.

Rämä, P., Sirri, L., & Serres, J. (2013). Development of lexical–semantic language system: N400 priming effect for spoken words in 18-and 24-month old children. *Brain And Language*, 125(1), 1-10.

Rescorla, L. (2002). Language and reading outcomes to age 9 in late-talking toddlers. *Journal of Speech, Language, and Hearing Research*, *45*(2), 360-371.

Rescorla, L. (2009). Age 17 language and reading outcomes in late-talking toddlers: Support for a dimensional perspective on language delay. *Journal of Speech, Language, and Hearing Research*, *52*(1), 16-30.

Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of verbal learning and verbal behavior*, *12*(1), 1-20.

Riordan, B., & Jones, M. N. (2011). Redundancy in perceptual and linguistic experience: Comparing feature-based and distributional models of semantic representation. *Topics in Cognitive Science*, *3*(2), 303-345.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge: MIT press.

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive psychology*, *7*(4), 573-605.

Rosch, E., Simpson, C., & Miller, R. S. (1976). Structural bases of typicality effects. *Journal of Experimental Psychology: Human perception and performance*, *2*(4), 491.

Rumelhart, D. E. (1990). Brain style computation: Learning and generalization. In F. Zornetzer, J. L. Davis, & C. Lau (Eds.), *An introduction to neural and electronic networks* (pp. 405–420). San Diego, CA: Academic Press Professional.

Smiley, S. S., & Brown, A. L. (1979). Conceptual preference for thematic or taxonomic relations: A nonmonotonic age trend from preschool to old age. *Journal of Experimental Child Psychology*, *28*(2), 249-257.

Smith, L. B. (2000). Learning how to learn words: An associative crane. *Becoming a word learner: A debate on lexical acquisition*, 51-80.

Smith, L. B. (2005). Shape: A developmental product. In L. Carlson, & van der Zee, E. (Eds.), *Functional features in language and space*, Oxford: Oxford University Press.

Smith, L. B., Colunga, E., & Yoshida, H. (2010). Knowledge as process: contextually cued attention and early word learning. *Cognitive science*, *34*(7), 1287-1314.

Smith, L. B., Jones, S. S., & Landau, B. (1996). Naming in young children: A dumb attentional mechanism? *Cognition*, 60(2), 143-171.

Smith, L. B., Jones, S. S., Landau, B., Gershkoff-Stowe, L., & Samuelson, L. (2002). Object name learning provides on-the-job training for attention. *Psychological Science, 13*(1), 13-19.

Smith, L. B., & Samuelson, L. (2006). An attentional learning account of the shape bias: Reply to Cimpian and Markman (2005) and Booth, Waxman, and Huang (2005). *Developmental Psychology*, 42, 1339–1343.

Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological review*, *81*(3), 214.

Smith, L. B., Suanda, S. H., & Yu, C. (2014). The unrealized promise of infant statistical word–referent learning. *Trends in cognitive sciences*, *18*(5), 251-258.

Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, *106*(3), 1558-1568.

Soja, N. N., Carey, S., & Spelke, E. S. (1991). Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. *Cognition*, *38*(2), 179-211.

Steyvers, M., & Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive science*, *29*(1), 41-78.

Storkel, H. L. (2001). Learning new words: Phonotactic probability in language development, *Journal of Speech, Language, and Hearing Research*, 44: 1321–37.

Storkel, H. L. (2009). Developmental differences in the effects of phonological, lexical and semantic variables on word learning by infants. *Journal of child language*, *36*(2), 291-321.

Storkel, H. L., & Adlof, S. M. (2009). The effect of semantic set size on word learning by preschool children. *Journal of Speech, Language, and Hearing Research*.

Storkel, H. L., Armbrüster, J., & Hogan, T. P. (2006). Differentiating phonotactic probability and neighborhood density in adult word learning. *Journal of Speech, Language, and Hearing Research*, *49*(6), 1175-1192.

Storkel, H. L., Bontempo, D. E., & Pak, N. S. (2014). Online learning from input versus offline memory evolution in adult word learning: Effects of neighborhood density and phonologically related practice. *Journal of speech, language, and hearing research*, *57*(5), 1708-1721.

Storkel, H. L., & Hoover, J. R. (2011). The influence of part-word phonotactic probability/neighborhood density on word learning by preschool children varying in expressive vocabulary. *Journal of Child Language*, *38*(3), 628-643.

Swingley, D., & Aslin, R. N. (2007). Lexical competition in young children's word learning. *Cognitive psychology*, *54*(2), 99-132.

Vigliocco, G., Vinson, D. P., Lewis, W., & Garrett, M. F. (2004). Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive psychology*, *48*(4), 422-488.

Vitevitch, M. S., Ercal, G., & Adagarla, B. (2011). Simulating retrieval from a highly clustered network: Implications for spoken word recognition. *Frontiers in psychology*, *2*, 369.

Waxman, S. R., & Hall, D. G. (1993). The development of a linkage between count nouns and object categories: Evidence from fifteen-to twenty-one-month-old infants. *Child development*, *64*(4), 1224-1241.

Woodward, A. L., Markman, E. M., & Fitzsimmons, C. M. (1994). Rapid word learning in 13- and 18-month-olds. *Developmental psychology*, *30*(4), 553.

Yoshida, H., & Smith, L. B. (2003). Shifting ontological boundaries: how Japanese-and English-speaking children generalize names for animals and artifacts. *Developmental Science*, *6*(1), 1-17.

Yu, C., & Smith, L. B. (2011). What you learn is what you see: using eye movements to study infant cross-situational word learning. *Developmental science*, *14*(2), 165-180.

Yurovsky, D., Bion, R. A. H., Smith, L. B., & Fernald, A. (2012). Mutual Exclusivity and Vocabulary Structure. In N. Miyake, D. Peebles & R. P. Cooper (Eds.), *Proceedings of the 34th annual meeting of the cognitive science society* (pp. 1197-1202). Austin Texas: Cognitive Science Society.

APPENDIX A. VALIDATING CORRECTED DEGREE

Here we conduct analyses designed to validate corrected degree as a metric that both 1) alleviates key problems with degree – namely, that degree correlates strongly with age and percentile, two individual difference measures of interest in this project – and 2) can be used to characterize items in ways that are useful for the primary research questions driving this work – specifically, we demonstrate that corrected degree more strongly relates to category density than degree. We conducted these analyses using 4,094 administrations (those that have both age and sex, so as to be able to assign a percentile) of the American English version of the MBCDI:WS from the Wordbank repository, using only the twelve items selected for study 1: airplane, bear, bicycle, boat, bug, bus, cat, chicken (the animal), duck, fish (the animal), train, and truck.

In this first set of analyses we first qualitatively and then quantitatively compare how degree versus corrected degree relate to age and percentile. Figure 19 shows degree and corrected degree as a function of age and percentile, for each of the twelve items of interest. Qualitatively, we note three main patterns of interest. First, and most importantly, while degree increases with both age and percentile, corrected degree remains relatively flat. Though corrected degree noticeably decreases with age for a couple words, the slopes are shallower than for degree. Second, the rank order of items for degree and corrected degree are generally the same. Finally, the variance in corrected degree for lower percentile participants is quite large, which aligns with findings that the vocabulary composition of lower percentile participants is much more variable than for higher percentile participants (e.g., Colunga & Sims, 2017; Perry & Samuelson, 2011). For our quantitative analyses, we created two linear mixed effects models (with administration and item as random effects) with age and percentile predicting degree and corrected degree (with all variable centered and scaled). The coefficients for the models are presented in Table 14.

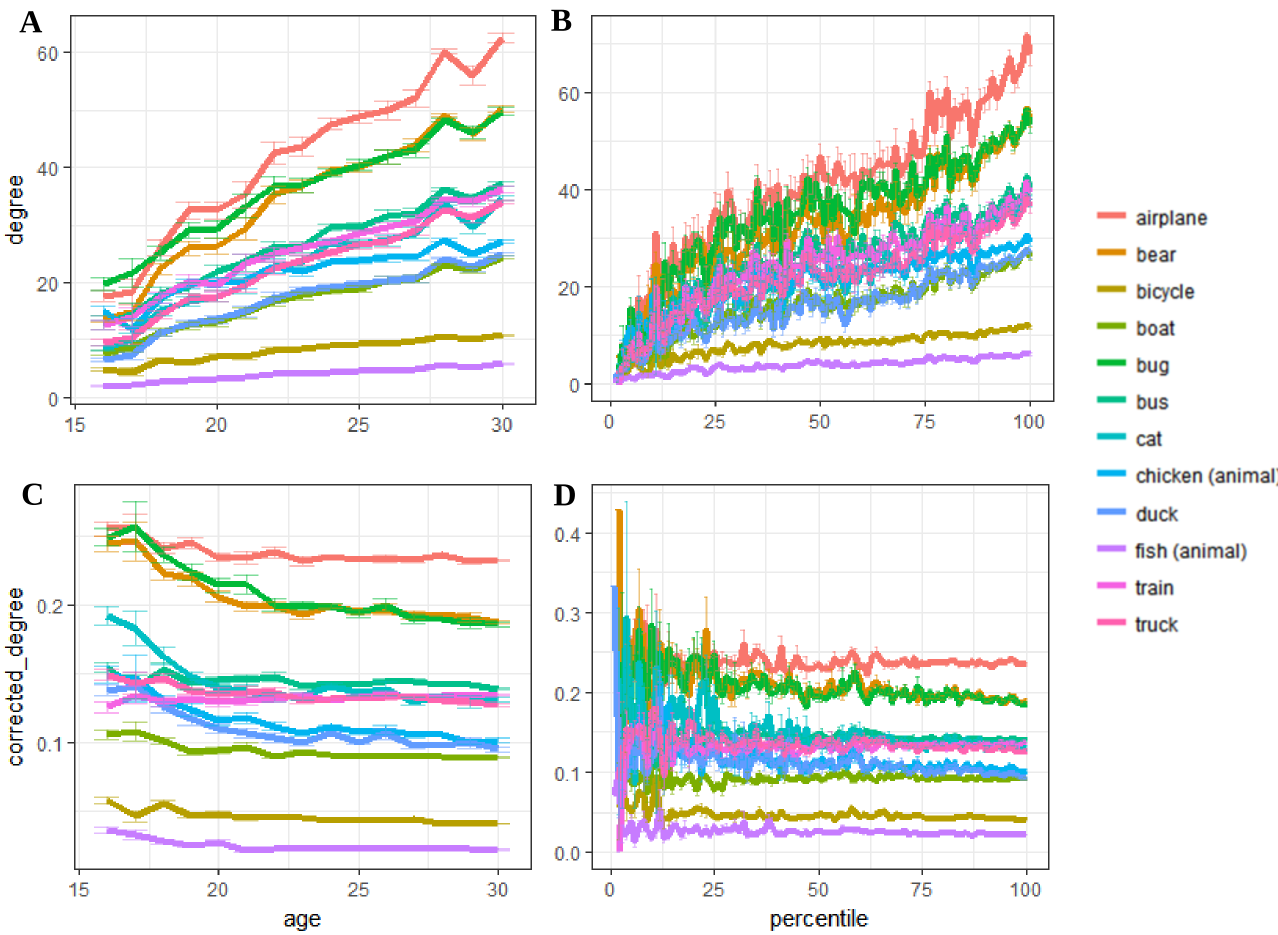

Figure 19. Figures of mean degree as a function (A) age and (B) percentile, and mean corrected degree as a function of (C) age and (D) percentile, for each of the twelve items used for the study in Chapter 2, using 4,094 administrations of the MBCDI:WS (with *SE* bars).

Table 14. Results for linear mixed effects models of age and percentile predicting degree and corrected degree

	Variable	Coef.	95% CI	t.val	p.val
Degree	Intercept	*0.335*	[-0.028, 0.699]	1.965	*.07*
	Age	**0.367**	[0.361, 0.373]	121.537	**<.001**
	Percentile	**0.349**	[0.343, 0.356]	105.652	**<.001**
Corrected Degree	Intercept	0.355	[-0.073, 0.783]	1.765	.10
	Age	**-0.111**	[-0.123, -0.096]	-14.61	**<.001**
	Percentile	**-0.086**	[-0.103, -0.07]	-10.28	**<.001**

Note. Coefficients are β_{std}.

The pattern of coefficients for the models supports the qualitative analyses. Most notably, the relations between the individual difference measures (age and percentile) and corrected degree are an order of magnitude weaker than for degree (per the t values).

We next explored whether degree and corrected degree predict the density of categories in individual toddlers' vocabularies. We calculated category density (for the twelve categories of nouns on the MBCDI:WS) as the number of items that an individual produces out of the total number of items for that category. We then created a linear mixed effects model (with administration and item as random effects) predicting category density as a function of degree and corrected degree. The coefficients for the model are presented in Table 15. Notably, corrected degree more strongly relates to category density than degree.

Table 15. Results for liner mixed effects model of category density as a function of degree and corrected degree.

Variable	Coef.	95% CI	t.val	p.val
Intercept	**-0.818**	[0.07, 0.844]	-22.95	**<.001**
Degree	**0.171**	[0.378, 0.39]	111.07	**<.001**
Corrected degree	**0.324**	[-0.101, 0.068]	167.85	**<.001**

Note. Coefficients are β_{std}.

Altogether, given our interest in the individual differences measures of age and percentile together with our desire to align with previous work using category density, the results presented here support the use of corrected degree for this project.

APPENDIX B. EXPERIMENTAL ITEMS

Here we present information about the items used for the empirical studies in chapters 2 and 3. Details regarding the known items used in Chapter 2 are presented in Table 16A and 16B, and details for the novel items used in Chapter 3 are presented in Table 17.

Table 16A. Item characteristics for known items used in Chapter 2.

Item	Group	VM & VFS Degree	Total Degree	Percep. NoF	VM & VFS NoF	CHILDES Freq.	Number Phonemes	Number Syllables	Bigram Freq.	Phon. ND	AoA	Proportion Produces
bus	high	33	50	10	6	6.84	3	1	7.95	7	20.44	0.82
boat	low	23	31	6	1	6.1	3	1	8.69	9	20.19	0.85
truck	high	30	32	5	4	7.87	4	1	6.63	4	18.56	0.93
bike	low	10	50	9	6	5.95	3	1	7.64	8	21.19	0.81
airplane	high	55	57	8	4	6.34	5	2	5.54	0	19.74	0.90
train	low	33	36	7	4	7.31	4	1	7.13	6	20.61	0.86
bug	high	43	44	6	5	6.37	3	1	7.95	19	20.95	0.83
chicken	low	23	72	10	6	6.91	5	2	6.37	1	22.31	0.76
bear	high	44	84	11	7	7.86	3	1	8.02	18	18.59	0.87
duck	low	22	57	11	7	8.01	3	1	7.31	13	16.44	0.94
cat	high	31	34	12	9	7.38	3	1	7.33	22	17.27	0.94
fish	low	5	71	5	2	7.43	3	1	6.92	3	18.44	0.91

Table 16B. Group-level statistics and comparisons of known items used in Chapter 2.

Measure	VM & VFS Degree	Total Degree	Percep. NoF	VFS & VM NoF	CHILDES Freq.	Number Phonemes	Number Syllables	Bigram Freq.	Phon. ND	AoA	Proportion Produces
High											
MN	39.33	57.33	10	6.17	1766.83	3.33	1.17	1679.07	12.33	19.01	0.89
SD	9.77	23.80	3.16	1.94	877.08	1.03	0.41	1105.31	9.89	1.11	0.05
Range	30-55	32-87	5-14	4-9	569-2621	2-5	1-2	254.93-3039.46	0-23	17.27-20-44	0.82-0.94
Low											
MN	19.33	48	7.5	4.17	1340.17	3.5	1.17	2125.07	6.67	19.64	0.87
SD	10.13	14.64	2.17	2.32	972.42	0.84	0.41	1922.36	4.32	1.85	0.05
Range	5-33	31-71	5-11	1-7	383-3005	3-5	1-2	966.16-5962.28	1-13	16.44-21.19	0.81-0.94
Comparison											
paired T-val	9.13	0.75	1.42	1.44	0.90	-0.28	0	-0.64	1.15	-0.94	0.62
p-value	**<.001**	.49	.22	.21	.41	.79	1	.55	.30	.39	.56

VM & VFS Degree = Item degree in networks built using on visual-motion and visual-form and surface features. Percep. NoF = Number of Perceptual Features. VFS & VM NoF = Number of Visual-Form and Surface and Visual Motion Features. Phon. ND = Phonological Neighborhood Distance (Colheart's N). AoA = Age of Acquisition. Proportion Produces = Proportion of 26-month-olds that produce the word on the MBCDI:WS.

Table 17. Stimuli characteristics for Novel Objects (List 1) used in Chapter 3.

Pairing	Label	Object	Feature	Feature Type	Degree (24mo/30mo)
1-high	Bosa	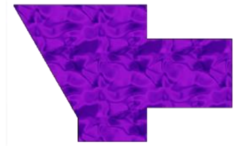	eats	visual-motion	11/15
1-low	Modi	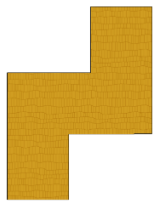	hops	visual-motion	2/2
2-high	Koba	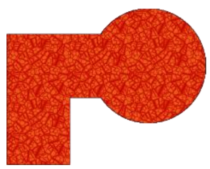	flies	visual-motion	8/9
2-low	Teebu	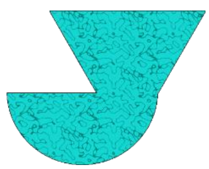	climbs trees	visual-motion	1/2
3-high	Coodle	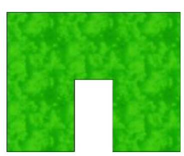	has wings	visual-form and surface	8/13
3-low	Tanzer	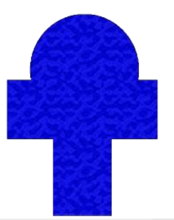	has fins	visual-form and surface	2/3

Object names were selected from the Novel Object and Unusual Name Database (Horst & Hout, 2014), and were controlled for: number of syllables, number of phonemes, and phonemic neighborhood density.

APPENDIX C. CHAPTER 2 PRE-REGISTRATION

Study Information

Title: Does Semantic Connectivity Influence Attention to and Subsequent Processing of Known Objects?

Authors: Ryan E Peters, Justin B Kueser, Arielle Borovsky

Affiliation: Purdue University

Rationale: Does semantic word knowledge influence the way toddlers attend to known objects and labels? A range of research provides evidence that toddlers' patterns of attention relate to their word learning outcomes (Colunga & Smith, 2005; Dixon & Smith, 2000; Kannass & Oakes, 2008; MacRoy-Higgins & Montemarano, 2016; Smith, 1995, 2000; Smith, Jones, & Landau, 1996; Smith, Jones, Landau, Gershkoff-Stowe, & Samuelson, 2002; Yu & Smith, 2011), and recent modeling work demonstrates how novel word learning and familiar word recognition may result from the same underlying mechanisms (McMurray, Horst, & Samuelson, 2012). Indeed, recent empirical work by Borovsky and colleagues provides evidence for similar facilitatory effects of lexico-semantic density in both novel word learning (Borovsky, Ellis, Evans, & Elman, 2015) and known word processing (Borovsky, Ellis, Evans, & Elman, 2016). However, less is known about more fine-grained influences of semantic word knowledge on attention to known objects. Here we ask: Does toddlers' attention to known objects vary with the number or quality of these objects' semantic relations with other words toddlers already know – or in other words, with the objects' patterns of semantic connectivity? Further, do these relationships affect subsequent object-label processing?

Research Question: The primary goal of this project is to explore how the semantic connectivity of known objects influences patterns of attention pre-labeling and subsequent object-label processing. A secondary goal, to be pursued in follow-up exploratory analyses, is to determine whether any semantic connectivity effects are related to individual differences such as age, attentional skill, word learning skill or lexico-semantic connectivity characteristics of the toddler's productive vocabulary.

To test these questions, 24- to 30-month-old toddlers will be exposed to known object-label pairs that vary in degree of semantic connectivity based on a set of feature production norms for early learned nouns (Peters, McRae, & Borovsky, in prep). We will measure pre- and post-labeling attention allocation/processing by measuring the toddlers' eye-fixations to object image pairs while listening to the label of one of the objects. In addition to the word learning task, we will use questionnaires to collect information from parents on individual differences of interest, including language background history, language environment, attentional skills, and productive vocabulary.

Hypotheses & predictions: We hypothesize that the semantic connectivity of known objects will influence (1) toddlers' patterns of attention and (2) subsequent patterns object-label processing. Evidence for this hypothesis comes from recent findings indicating that patterns of known word processing relate to the density of the semantic category of a word in the context of the individual toddler's vocabulary (Borovsky, Ellis, Evans, & Elman, 2016).

We consider four possible outcomes, each of which have different implications for mechanisms underlying how attention and connectivity support learning:

1. Attentional biases drive looks to objects with high connectivity, and support subsequent processing of labels for those objects.

2. Attentional biases drive looks to objects with low connectivity, and support subsequent processing of labels for those objects.

3. High semantic connectivity supports processing of object-labels, irrespective of attentional biases.

4. Low semantic connectivity supports processing of object-labels, irrespective of attentional biases.

Sampling Plan

Recruitment procedures: 24- to 30-month-olds will be recruited via (1) a participant registry run by the department of Speech, Hearing, and Language Sciences at Purdue University and (2) a database of infants and children from 2 months to 4 years of age.

Inclusionary and exclusionary criteria: During initial recruitment, families will be asked preliminary screening questions to determine if the toddlers fit the following inclusionary criteria:

1. Normal vision (based on parent report).

2. Normal hearing (based on parent report) and no history of chronic ear-infections.

3. No reported history of neurological or cognitive impairment.

4. Normal birth history: were not born preterm (<37weeks) and low birth weight (< 5lb8oz).

5. Not exposed to language other than English for more than 8 hours a week (~10% of waking hours).

Sample size: 60 toddlers who meet inclusionary and exclusionary criteria will be tested

Sample size rationale: Power analyses indicated that for a multiple regression with 4 parameters, to achieve a power of .8 (enough to detect a small to moderate effect size with *Cohen's f2*=.25) for a significance level of .05, a sample size of 48 is necessary. Thus, assuming at attrition rate of ~%20 in the eye-tracking portion of the experiment (due to fussiness, etc.), a sample size of ~60 is necessary.

Design Plan

Study type: Experiment

Blinding: Research personnel who interact directly with the study subjects will not be aware of the assigned treatments. Additionally, the experimental protocol is administered via computer using a standardized method for all children.

Study design: We have a within-subjects design with 1 factor (semantic connectivity) that can be treated as either continuous or categorical (with 2 levels). The semantic connectivity (high/low) will be counterbalanced.

Experiment structure: The trials of interest are filler trials for a larger word-learning study. There are three blocks in the experimental task, with 8 trials per block, for a total of 24 trials. The trials have four time periods (following precedent for known word processing, e.g., Borovsky et al., 2016): (1) *Preview Period*: first 1500 ms of the trial (pre-audio stimulus), (2) *Test Period*: from 300 to 4000 ms post audio stimulus onset, (3*) Early Test Period*: from 300 to 1500[2] ms post audio stimulus onset, and (4) *Late Test Period*: from 1500 to 4000 ms post audio stimulus onset.

[2] An error was made in the original pre-registration document. The Early Test Period was described as going from 300 to 1300 ms post audio stimulus onset, and the Late Test Period as going from 1300 to 4000 ms. This description was mistakenly carried over from the pre-registration for a separate study – exploring online processing of novel words – that this data was collected together with, and for which the pre-registration documents were written in parallel. This correction was made prior to the finalization of the dataset.

In addition to the experimental task, we will use the following questionnaires to collect information from parents on individual differences of interest:

Background Questionnaire: This questionnaire is designed to collect a range of background information, including birth history, demographic information, language history, and information on the language learning environment.

Early Childhood Behavior Questionnaire (ECBQ: Putnam, Gartstein, & Rothbart, 2006): This questionnaire asks about a range of temperamental characteristics of the child, however our focus will be on the sections that collect information on attentional skills.

MacArthur-Bates Communicative Developmental Inventory (MBCDI; Fenson et al., 2007): This is a parental questionnaire of early-produced words. We will use this to estimate vocabulary size, connectivity patterns, and individualized estimates of semantic connectivity for the novel objects.

Experimental Word Knowledge Questionnaire: This questionnaire contains all of the words in the experiment. Parents rate both comprehension and production on scales from 1 (definitely does not understand/say) to 4 (definitely understands/says).

Randomization: Participants will be randomly assigned to one of 6 lists, which have pseudo-randomized order of presentation.

Variables

Manipulated variables:

IV 1A: Semantic *Connectivity (categorical)*. Pairs of objects consist of one *high-connectivity* and one *low-connectivity* object. High- and low-connectivity objects have perceptual semantic features that are relatively common or rare, respectively, in the known concepts of children in the target age-range (per a set of feature production norms for early learned nouns; Peters, McRae, &

Borovsky, in press). High- and low-connectivity objects share perceptual features with, on average, 48 and 13.5 concepts in a normative 30-month-old's productive vocabulary, respectively. We focused on perceptual features, based on evidence that they play a particularly important role in early lexico-semantic development (Peters & Borovsky, in revision).

IV 1B: Semantic *Connectivity (continuous).* An individualized, continuous measure of semantic connectivity will be calculated for each object, for each child. This measure will be calculated as the number of concepts in the child's productive vocabulary that share features with the object.

Measured variables: For DVs 1-2, we will create metrics that characterize the bias towards one object versus another, called log gaze-proportion ratio. These metrics are calculated as the log of the proportion of looks to the object of interest over the proportion of looks to the distractor.

DV 1: *Pre-labeling Bias*. Measure of attention allocation pre-labeling (Log gaze-proportion ratio during preview period).

DV 2: *Accuracy*. Measure of attention allocation post-labeling (Log gaze-proportion ratio for the target, post-labeling in RRT-REC test period). Based on pilot testing, there will be 3 time-windows of interest: (1) Test period: 300 ms post audio stimuli onset to end of trial. (2) Early test period: the first 1000 ms of the test period. (3) Late test period: the remainder of test period after the early fixation period.

Analysis Plan

Statistical models:

Analyses exploring whether biases relate to semantic connectivity.

1a. Paired two-way t-test comparing participant level *Pre-Labeling Bias* (DV 1) between high (1) and low (0) *Semantic Connectivity* (*categorical*; IV 1A) words. Higher pre-labeling bias values for either high or low connectivity objects would support the hypotheses that attentional biases drive looks to objects with high or low semantic connectivity, respectively.

1b. Linear mixed effects model with random effects for participants and items, with the Dependent Variable: trial level *Pre-Labeling Bias* (DV 1), and Independent Variable: *Semantic Connectivity* (*continuous;* IV 1B) of objects. All variables will be scaled and centered. A positive or negative relation between pre-labeling bias and semantic connectivity would support the hypotheses that attentional biases drive looks to objects with high or low semantic connectivity, respectively.

Analyses exploring whether object-label processing relates to semantic connectivity.

2a. Paired one-way (positive) t-test comparing participant level *Accuracy* (DV 2) between high (1) and low (0) *Semantic Connectivity* (*categorical*; IV 1A) objects. We will run three versions of this test, one for each of the pre-determined test windows for Accuracy. Higher accuracy values for either high or low connectivity objects would support the hypotheses that object-label processing is supported by high or low levels of semantic connectivity, respectively.

2b. Linear mixed effects model with random effects for participants and items, with the Dependent Variable: trial level *Accuracy* (DV 2), and Independent Variables: *Semantic Connectivity* (*continuous;* IV 1B) of objects, *Pre-labeling Attention* (DV 1), and the *interaction* between Semantic Connectivity and Pre-labeling Attention (IV 1B * DV 1). We will run three versions of this model, one for each of the pre-determined test windows for Accuracy. All variables will be scaled and centered. A positive or negative relation between accuracy and semantic connectivity would support the hypotheses that object-label processing is supported by high or low levels of semantic connectivity, respectively. A positive relation between accuracy and pre-labeling bias would indicate that pre-labeling attention supports object-label processing.

Finally, a significant relation between the accuracy and interaction terms would indicate that semantic connectivity plays a role in pre-labeling biases to objects which then supports subsequent processing of labels for those objects.

Exploratory follow-up analyses of individual differences: We will explore relations between the main effects described above and individual differences such as age, attentional skill, word learning skill and lexico-semantic connectivity characteristics of the toddler's vocabulary.

Inference criteria: We will use two methods. First, we will use the standard $p < .05$ criteria (determined via bootstrapped confidence intervals for the linear mixed effects models) for determining if the linear mixed effects models suggest that the results are significantly different from those expected if the null hypothesis were correct. Second, for marginal effects with $p < .1$, we will use the $d >= |.35|$ (i.e. small/medium effect size) criteria for determining if results can be considered when making inferences.

Data exclusion: We will exclude any trials from analyses that correspond to words for which parents mark comprehension as 1 or 2, on the Experimental Word Knowledge Questionnaire. Individual trials in the RRT will be removed if the percentage of track-loss exceeds 80 percent in the Test Period, following Borovsky et al., 2016. Finally, the data from individual participants will be removed from the RRT if, after the removal of individual trials based on track-loss and the Experimental Word Knowledge Questionnaire, they do not have at least 2 high-connectivity and 2 low-connectivity RRT-RET data points.

Missing data: If a participant does not complete at least half of the eye-tracking task or their parent does not complete the vocabulary checklist, that subject will not be included in the analyses.

APPENDIX D. CHAPTER 3 PRE-REGISTRATION

Study Information

Title: Does Semantic Connectivity Influence Attention to Novel Objects and the Learning of Labels?

Authors: Ryan E Peters, Justin B Kueser, Arielle Borovsky

Affiliation: Purdue University

Rationale: Does semantic word knowledge influence the way toddlers attend to novel objects and learn labels? A range of research provides evidence that toddlers' patterns of attention relate to their word learning outcomes. This research has explored attention at multiple levels of analysis, including: participant level measures of temperamental attentional characteristics (Dixon & Shore, 1997; Dixon & Smith, 2000), task level measures of attention in a word learning context (Kannass & Oakes, 2008), and stimuli level measures exploring relations between attention allocated to specific stimuli and subsequent learning outcomes (MacRoy-Higgins & Montemarano, 2016; Yu & Smith, 2011). Furthermore, a rich line of research has explored how attentional biases towards specific object features, in particular shape, develop over the course of early word learning and subsequently facilitate it (e.g., Colunga & Smith, 2005; Smith, 1995, 2000; Smith, Jones, & Landau, 1996; Smith, Jones, Landau, Gershkoff-Stowe, & Samuelson, 2002). This research has revealed broad relations between word knowledge and attention – for example showing that the tendency to attend to shape only emerges after toddlers know a large number of nouns (e.g., Samuelson & Smith, 1999). However, less is known about more fine-grained effects of semantic knowledge on attention to novel objects. Here we ask: Does toddlers' attention to novel objects vary with the number or quality of these objects' semantic relations with the words toddlers already know – or in other words, with the objects' patterns of semantic connectivity? Further, do these relationships affect subsequent novel word learning outcomes?

Research Question: The primary goal of this project is to explore how the semantic connectivity of novel objects influences patterns of attention and subsequent object-label learning outcomes. A secondary goal, to be pursued in follow-up exploratory analyses, is to determine whether any semantic connectivity effects are related to individual differences such as age, attentional skill, word learning skill or lexico-semantic connectivity characteristics of the toddler's productive vocabulary.

To explore these issues, 24- to 30-month-old toddlers will be exposed to novel objects designed to have semantic features that are either common or rare in the known concepts of children in the target age-range (according to a set of feature production norms for early learned nouns; Peters, McRae, & Borovsky, in press). Next, each novel object will be named several times. Finally, in the test task, we determine the degree of word learning by measuring the toddlers' eye-fixations to pairs of object images while they listen to the label of one of the objects. In addition to the word learning task, we will use questionnaires to collect information from parents on individual differences of interest, including language background history, language environment, attentional skills, and productive vocabulary.

Hypotheses: We hypothesize that the semantic connectivity of novel objects will influence (1) toddlers' patterns of attention and (2) subsequent object-label learning outcomes. Evidence for this hypothesis comes from recent findings indicating that word learning skill relates to the semantic structure of a toddler's vocabulary (Beckage, Smith, & Hills, 2011; Borovsky, Ellis, Evans, & Elman, 2016). One explanation for these findings is that attentional biases, which aid in word learning, depend on semantic relations between to-be-learned words and known vocabulary.

We consider four possible outcomes, each of which have different implications for mechanisms underlying how attention and connectivity support learning:

1. Attentional biases drive looks to novel objects with high connectivity features, and support subsequent learning of object-label associations for those objects.

2. Attentional biases drive looks to novel objects with low connectivity features, and support subsequent learning of object-label associations for those objects.

3. High semantic feature connectivity supports the learning of object-label associations, irrespective of attentional biases.

4. Low semantic feature connectivity supports learning of object-label associations, irrespective of attentional biases.

Prior work has demonstrated that if children know a large fraction of concepts in the category of which a novel object belongs (e.g., vehicles) learning is facilitated for that novel object (Borovsky et al., 2016). Concepts in categories typically share many features, meaning this result may be due to toddlers noting overlap between features in the novel objects and many of the concepts they know, and providing support for outcomes 1 and 3.

Sampling Plan

Recruitment procedures: 24- to 30-month-olds will be recruited via (1) a participant registry run by the department of Speech, Hearing, and Language Sciences at Purdue University and (2) a database of infants and children from 2 months to 4 years of age.

Inclusionary and exclusionary criteria: During initial recruitment, families will be asked preliminary screening questions to determine if the toddlers fit the following inclusionary criteria:

1. Normal vision (based on parent report).
2. Normal hearing (based on parent report) and no history of chronic ear-infections.
3. No reported history of neurological or cognitive impairment.
4. Normal birth history: were not born preterm (<37weeks) and low birth weight (< 5lb8oz).
5. Not exposed to language other than English for more than 8 hours a week (~10% of waking hours).

Sample size: 60 toddlers who meet inclusionary and exclusionary criteria will be tested

Sample size rationale: Power analyses indicated that for a multiple regression with 4 parameters, to achieve a power of .8 (enough to detect a small to moderate effect size with *Cohen's f2*=.25) for a significance level of .05, a sample size of 48 is necessary. Thus, assuming at attrition rate of ~%20 in the eye-tracking portion of the experiment (due to fussiness, etc.), a sample size of ~60 is necessary.

Design Plan

Study type: Experiment

Blinding: Research personnel who interact directly with the study subjects will not be aware of the assigned treatments. Additionally, the experimental protocol is administered via computer using a standardized method for all children.

Study design: We have a within-subjects design with 1 factor (semantic connectivity) that can be treated as either continuous or categorical (with 2 levels). The semantic connectivity (high/low) will be counterbalanced.

Experiment structure: The experimental task consists of three blocks, each with three sub-tasks:

Feature Exposure Task (FET): In the feature exposure task participants will first be presented two novel objects side by side to determine any initial visual bias (Visual Bias). Next, they will see videos demonstrating semantic features of the novel objects (Feature Exposure). There will be one video for each object, each presented twice. Finally, they will once again be presented with the two novel objects side by side to determine any bias that results from the feature exposure (Feature Bias).

Naming Exposure Task (NET): In the naming exposure task participants will listen to each novel object being labeled with a novel name (2 trials for each object).

Retention and Recognition Task (RRT): In the retention and recognition task participants will see a pair of objects side by side, one of which will be labeled with an audio stimulus. There will be four trials testing retention of the novel objects (RET), and eight filler trials containing pairs of known objects. The RET trials have four time periods (determined via pilot testing): (1) *Post-naming Bias Period*: first 1500 ms of the trial (pre-audio stimulus), (2) *Test Period*: from 300 to 4000 ms post audio stimulus onset, (3) *Early Test Period*: from 300 to 1300 ms post audio stimulus onset, and (4) *Late Test Period*: from 1300 to 4000 ms post audio stimulus onset.

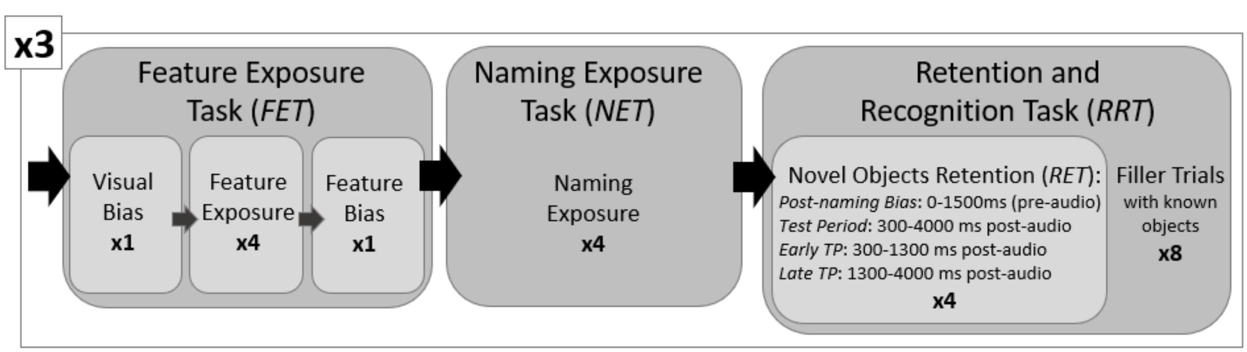

Figure D-1. A diagram of one block of the experiment.

In addition to the experimental task, we will use the following questionnaires to collect information from parents on individual differences of interest:

Background Questionnaire: This questionnaire is designed to collect a range of background information, including birth history, demographic information, language history, and information on the language learning environment.

Early Childhood Behavior Questionnaire (ECBQ: Putnam, Gartstein, & Rothbart, 2006): This questionnaire asks about a range of temperamental characteristics of the child, however our focus will be on the sections that collect information on attentional skills.

MacArthur-Bates Communicative Developmental Inventory (MBCDI; Fenson et al., 2007): This is a parental questionnaire of early-produced words. We will use this to estimate vocabulary size, connectivity patterns, and individualized estimates of semantic connectivity for the novel objects.

Experimental Word Knowledge Questionnaire: This questionnaire contains all of the words in the experiment, including words in the descriptions of the features (e.g., “fly”, “hop”, “fin”). Parents rate both comprehension and production on scales from 1 (definitely does not understand/say) to 4 (definitely understands/says).

Randomization: Participants will be randomly pre-assigned to one of 6 lists, which have pseudo-randomized order of presentation.

Variables

Manipulated variables:

IV 1A: Semantic *Connectivity (categorical)*. Pairs of novel objects consist of one *high-connectivity* and one *low-connectivity* object. High- and low-connectivity objects are designed to have semantic features that are common or rare, respectively, in concepts in the productive vocabulary of children in the target age-range (per a set of feature production norms for early learned nouns; Peters, McRae, & Borovsky, in press). High- and low-connectivity objects have features that are, on average, in 9 and 1.67 words in a normative 24-month-old's productive vocabulary, respectively, and 12.33 and 2.33 words in a normative 30-month-old's productive vocabulary.

IV 1B: Semantic *Connectivity (continuous).* An individualized, continuous measure of semantic connectivity will be calculated for each object, for each child. This measure will be calculated as the number of concepts on the MBCDI that are in the child's productive vocabulary that share the feature assigned to the novel object.

Measured variables: For DVs 1-4, we will create metrics that characterize the bias towards one object versus another, called <u>log gaze-proportion ratio</u>. These metrics are calculated as the log of the proportion of looks to the object of interest over the proportion of looks to the distractor.

DV 1: *Visual Bias*. Measure of attention allocation before novel object is assigned a feature (Log gaze-proportion ratio during FET: Visual Bias).

DV 2: *Feature Bias*. Measure of attention allocation after novel object is assigned a feature (Log gaze-proportion ratio during FET: Feature Bias).

DV 3: *Post-naming Bias*. Measure of attention allocation pre-labeling in the RRT-RET (Log gaze-proportion ratio during RRT-RET pre-audio period).

DV 4: *RET Accuracy*. Measure of attention allocation post-labeling (Log gaze-proportion ratio for the target, post-labeling in RRT-RET test period). There will be 3 time-windows of interest, described in Design Plan.

Analysis Plan

Statistical models:

Analyses exploring whether biases relate to semantic connectivity.

1a. Paired two-way t-test comparing participant level *Feature Bias* (DV 2) between high (1) and low (0) *Semantic Connectivity (categorical;* IV 1A) novel objects. Higher feature bias values for either high or low connectivity objects would support the hypotheses that attentional biases drive looks to novel objects with high or low connectivity, respectively.

1b. Linear mixed-effects model with random effects for participants and items, with the Dependent Variable: trial level *Feature Bias* (DV 2), and Independent Variables: *Semantic Connectivity (continuous)* of novel words (IV 1B) and *Visual Bias* (DV 1). All variables will be scaled and centered. A positive or negative relation between feature bias and semantic connectivity would support the hypotheses that attentional biases drive looks to novel objects with high or low semantic connectivity, respectively, while considering initial visual biases.

Analyses exploring whether object-label learning outcomes relate to semantic connectivity.

2a. Paired one-way (positive) t-test comparing participant level *RET Accuracy* (DV 4) between high (1) and low (0) *Semantic Connectivity* (*categorical*; IV 1A) novel objects. We will run three versions of this test, one for each of the pre-determined test windows for RET Accuracy. Higher accuracy values for either high or low connectivity novel objects would support the hypotheses that object-label learning is supported by high or low levels of connectivity, respectively.

2b. Linear mixed effects model with random effects for participants and items, with the Dependent Variable: trial level *RET Accuracy* (DV 4), and Independent Variables: *Semantic Connectivity (continuous)* of novel words (IV 1B), *Visual Bias* (DV 5), *Feature Bias* (DV 6), and *Post-naming Bias* (DV 7). All variables will be scaled and centered. We will run three versions of this model, one for each of the pre-determined test windows for RET Accuracy. A positive or negative relation between accuracy and semantic connectivity would support the hypotheses that object-label learning is supported by high or low levels of semantic connectivity, respectively, while considering visual, feature, and post-naming biases.

Exploratory follow-analyses of individual differences: We will explore relations between the main effects described above and individual differences such as age, attentional skill, word learning skill and lexico-semantic connectivity characteristics of the toddler's vocabulary.

Inference criteria: We will use two methods. First, we will use the standard $p < .05$ criteria (determined via bootstrapped confidence intervals for the linear mixed effects models) for determining if the linear mixed effects models suggest that the results are significantly different from those expected if the null hypothesis were correct. Second, for marginal effects with $p < .1$, we will use the $d >= |.35|$ (i.e. small/medium effect size) criteria for determining if results can be considered when making inferences.

Data exclusion: We will exclude any trials from analyses that correspond to words for which parents mark comprehension as 1 or 2, on the Experimental Word Knowledge Questionnaire. Individual trials in the RRT will be removed if the percentage of track-loss exceeds 80 percent in the Test Period, following Borovsky et al., 2016. Finally, the data from individual participants will be removed from the RRT if, after the removal of individual trials based on track-loss and the Experimental Word Knowledge Questionnaire, they do not have at least 2 high-connectivity and 2 low-connectivity RRT-RET data points.

Missing data: If a participant does not complete at least half of the eye-tracking task or their parent does not complete the vocabulary checklist, that subject will not be included in the analyses.