DO OLD OBJECTS HELP INFANTS PAY ATTENTION TO NEW ONES?

by

ALLISON M. ZHOU

A THESIS

Presented to the Department of Biology
and the Robert D. Clark Honors College
in partial fulfillment of the requirements for the degree of
Bachelor of Science

June 2018

An Abstract of the Thesis of

Allison M. Zhou for the degree of Bachelor of Science
in the Department of Biology to be taken June 2018

Title: Do Old Objects Help Infants Pay Attention to New Ones?

Approved: ____

Dr. Caitlin Fausey

Babies' first words are often names for objects that appear frequently in their lives; could these frequent objects also help babies learn the names of other, less common, objects, too? Word learning is a multifaceted developmental process that involves attention, memory, and generalization. In the present study we focus on attention. We take inspiration from the observation that statistical non-uniformity governs many visual and auditory aspects of the world; the images we see and the words we hear are largely structured so that there are a small number of highly ubiquitous items and a large number that are much less prevalent. Here, 16- to 30-month-old infants observe pictures of novel objects sampled from a uniform distribution where they see each object an equal number of times, or from a non-uniform distribution where they see one object five times more often than others. We predict that infants will pay more attention to sequences of objects sampled from a non-uniform distribution, which has higher rates of repetition and allows learners to compare newly seen objects to a more familiar one. This method begins to address the challenge of how to precisely examine consequences of patterns of input in infants' everyday lives by bringing such patterns into the lab. This method also moves beyond business-as-usual

by explicitly testing a hypothesized sweetspot between repetition and novelty, which could be a generally fruitful approach used to reduce 'fuss out' rates in infant behavioral studies. Sustained attention gives infants the opportunity to encounter and learn about more objects. Our research will yield new insight into infant attention in the context of word learning.

Acknowledgements

I would like to thank Professor Caitlin Fausey for not only all of her help with my thesis, but also for her willingness and eagerness to teach, guide, and support me in my academic and professional endeavors throughout the last three years. I would also like to thank my Thesis Committee, which includes Professor Fausey, as well as Professor Louise Bishop and Professor Terry Takahashi. I would also like to thank the Learning lab team for their contributions. A special thank you to my fellow lab member and dear friend, Vinitha Gadiraju, for her help with coding. Lastly, thank you to all of my professors, friends, and family members for their encouragement and support throughout the thesis process.

Table of Contents

List of Figures

List of Tables

Introduction

At every turn the physical world around us offers a multitude of learning opportunities. Our everyday visual, auditory, and sensory environments are datasets with many statistical regularities. Evidence shows that human infants are attuned to these statistics in several domains including, but not limited to, vision, language, and music (Saffran & Aslin, 2002; Saffran, Aslin & Newport, 1999). Within this field of developmental psychology, there is a focus on characterizing infants' statistical learning abilities. What regularities are available in everyday life and which ones babies attend to are directly linked to what, when, and how babies learn. Infants make strong developmental strides within their first two years of life; visual object recognition and object naming are just a couple of the important skills babies grow to develop. Infants are strong statistical learners and are able to pick up quite a lot of information from their experiences and environments.

Many developmental studies specifically seek to characterize how infants attend to and learn from information presented in structured distributions. In-lab studies show that stimulus repetition, an aspect of statistical distribution, is helpful for learning words over time (Mather & Plunkett, 2009). We also know that variability helps with generalization and word learning; specifically, seeing a variety of objects within a category leads infants to better abstract both individual and global category organization *and* to learn more outside of the laboratory (e.g., Perry, Samuelson, Malloy & Schiffer, 2010). Many of these types of studies use training sets of objects and their names to present and/or teach to infants in the lab. Oftentimes, learning assessments are given to infants to test if infants learn to link names with objects, generalize to new instances, and make novel noun generalizations (NNGs), as well as overall vocabulary acquisition outside of the lab (Perry et al., 2010). Variable exemplars and repetition together have been shown to help facilitate word learning in babies; there is strong evidence that providing young children with the opportunity to compare across multiple exemplars during word learning tasks facilitates retention of recently learned words (Twomey, Ranson & Horst, 2013).

These aforementioned studies are a representative selection of the research done on early word learning and object naming. These studies typically use *uniform* distributions for in-lab experimental designs where specific word instances and object types are repeated and presented an equal number of times. For example, in a standard object naming and learning study comparing three different novel object types (e.g., "dax", "modi", and "gasser"), each object would usually be presented an equal number of times to the infant. This is a valid methodological choice for many reasons depending on the research question. While in-lab developmental studies arguably create an inherently superficial environment, if they seek to understand how babies learn “in the wild” they ought to strive to model real-life infant experiences as closely as possible. That is, in-lab studies could aim to understand the learning problems that infants actually face and solve (Smith et al., 2018; see Lee, Cole, Golenia & Adolph, 2017 for analogous arguments about motor development). This then leads to the question: what are some of the statistical regularities in the real-world? And what type of structure is available in infants’ everyday environments? Are distributions found “in the wild” largely uniform?

Infants encounter objects in highly non-uniform distributions in the wild.

Recent advances in characterizing structure in infants' environments reveal that real world distributions are largely *not* uniform like the training sets used in many of the aforementioned in-lab developmental studies. In fact, strong statistical *non-uniformity* governs many visual and auditory aspects of the world; the images we see and the words we hear are largely structured in a way so that there are a small number of highly ubiquitous items and a large number that are much less prevalent. For example, one analysis of an annotated distribution of objects in a corpus of Internet searched images showed that 9 objects out of 200 account for 50% of all image data (Salakhutdinov, Torralba & Tenenbaum, 2011, Table 1). Even the way we speak is highly skewed towards certain types of constructions and verb-types (e.g., Casenhiser & Goldberg, 2005, Table 1; see also Piantadosi, 2014). Most relevant to infants, head-camera data shows that young infants spend about 25% of their day looking at faces, among which are a few faces that appear with very high frequency (Jayaraman, Fausey & Smith, 2015, Table 1). Within specific activities like mealtimes, the set of objects across infant-views are also extremely skewed with a handful of objects appearing much more frequently than others (Clerkin et al., 2016). The mismatch between the statistical structure of faces and objects in infants' everyday experiences (non-uniform) and training sets used in developmental modeling and in-lab experiments on statistical learning (largely uniform) is striking. Because these distributions are characterized by different degrees of repetition and variability, and we know that both repetition and variability impact learning, it is worth exploring potential consequences that different distributions have for infant attention and learning. The present study proposes to implement a non-uniform object distribution in the lab and to pit it against the standard uniform distribution.

Domain	**Example type**	**Non-uniform finding**
Language	Language structure: Better lexical learning with non-uniformity in language Wonnacott, Brown, & Nation, 2017	6-year-olds and adults show stronger lexical learning when given skewed language examples where a majority particle co-occurrs with most nouns.
	Verb use: Specific verbs are used much more often in natural language of mothers to young children Casenhiser & Goldberg, 2005	"put" is the verb in ~40% of the instances in specific phrases in mothers' speech. e.g. Mom "puts" the book on the table. ("Puts" is used most frequently over less frequent verbs like "places" or "lays").
Visual Objects	Objects in outdoor images: Large object database showing object counts in outdoor and indoor images Spain & Perona, 2007	Windows are the most frequently appearing object in outdoor image examples with 2-10X the frequency of lower count objects such as a locker or coffin.
	Objects in Google images: Salakhutdinov, Torralba & Tenenbaum, 2011	9 objects out of 200 account for 50% of all image training data.
	Faces available to infants: Skewed distribution of faces in infant views Jayaraman, Fausey, & Smith, 2015	Infants spend about 25% of their day looking at faces, among which are a few high frequency individuals.

Infant Egocentric Views	Meal-time objects: Headcamera data of infant at-home mealtime events Clerkin, Hart, Rehg, Yu & Smith, 2016	Seven items of the 745 unique object types account for 33% of all reported object instances.

Table 1: Motivation for non-uniform distribution from a review of papers showing non-uniform distributions in real-world instances.

The pervasive nature of non-uniform distributions of objects, faces, and language use are recently well documented (see Table 1), but what are the learning implications of these non-uniform statistics? Studies show that infants are attuned to patterns in their environment and can actually extract structure from their linguistic and visual environments (Saffran & Aslin, 2002). While statistical learning contributes to helping infants learn to speak, exactly how these everyday repeated events capture infant attention and how they contribute to early learning is unknown. Word learning is a multifaceted developmental process that involves attention, memory, and generalization mechanisms. In the present study we focus on attention.

Infant Attention & Vision

Developing one's attention is an important foundational skill to word learning. Furthermore, an ability to pay attention for longer amounts of time and to suppress distractions consequently means more focused learning opportunities. Research shows that when infants learn object categories by attending to the shape of objects, they are better able to generalize this information to learn new object categories and also show a dramatic increase in acquisition of new object names outside of the laboratory as well (Smith et al., 2002). The question of what captures infant attention is an interesting research endeavor. We know that if a sequence is too simple or contrastingly too complex, infants will lose attention and look away faster (Kidd & Hayden, 2015). Perhaps a specific number of objects types or level of repetition and variability is more interesting and attention-grabbing. Perhaps an "attentional sweet-spot" exists where certain visual sequences are better able to hold infant attention based on their statistical structure.

The everyday visual environment represents a sea of information that far exceeds that which an infant can process in any given moment. Capacity limits within the retina and cerebral cortex make it impossible to visually process everything at once, but the visual system can select a subset of available visual information to attend to (Castelhano, Mack & Henderson, 2009). A critical objective in the study of developmental behavior is to understand what influences and captures infant sustained attention. One form of sustained attention is *selective* sustained attention, which includes the ability to maintain focus on a specific object or task such as viewing sequences of sensory input for an extended period of time (Fisher & Kloos, 2016). Here we examine if the structure of the environment, in the form of different object distributions, may support different degrees of sustained attention. Attention is sensitive to both content and context and so we hypothesize that infant attentional differences will arise when infants sample novel objects from a uniform vs. non-uniform distribution in the present study.

Infant Experiences: A concrete example

Think for a moment about what experiences with bowls an infant may have before learning the word "bowl". Perhaps the baby always eats out of a small, blue bowl, while her family owns two other sets of red bowls and wooden bowls. These are all bowl object-types the infant will encounter, but her own small, blue bowl is likely the one she sees and eats from the most often. Occasionally, the baby may see rare types of bowls like a crystal bowl or large metal mixing bowl. These experiences seeing many different types of bowls accumulate and over time the infant will likely see a large range of bowl-types. She will encounter most of the bowl-types infrequently, while her amount of experience with her own small, blue bowl stays consistently and relatively high. Likely, the infant becomes most familiar with her own bowl because she routinely sees it during mealtimes. By hypothesis, the large range of bowl-types an infant sees and the repetition of certain bowls all contribute to helping her learn and generalize an idea of what a "bowl" is.

A key factor to learning a new word is repetition, which has been shown to help with learning (Horst, Parsons & Bryan, 2011). Repetition is beneficial for learning in many ways. There is evidence that repeated exposure and visual familiarity with objects and words is helpful on both a short and long timescale (Clerkin et al., 2016; Schwab & Lew-Williams, 2016). Experiments on novel word learning with infants in the lab indicate that visual familiarity with objects prior to their naming specifically helps enhance learning and retention of the link between name and object (Fennel, 2011). Visual familiarity and repetition support perceptual and memory processes and may be critical to statistical learning (Clerkin et al., 2016). The recently articulated 'pervasiveness hypothesis' suggests that it may be helpful for an infant to see one object type the most (e.g., her own small, blue bowl) and this object, which is repeated most often, can act as a perceptual anchor (Clerkin et al., 2016). We know that repetition is important to learning, but this raises the question: what type or level of repetition is best at capturing infant attention? Research shows that infant attention often follows a "goldilocks effect" with infants seeking out intermediate rates and levels of information. Infants prefer to allocate their attention to sequences of intermediate difficulty perhaps to spend cognitive resources on events that have the most potential for understanding and learning (e.g., Kidd & Piantadosi, 2012; Gerken & Balcomb, 2012; see also Hunter & Ames, 1988).

Sequential Attention

The sequential attention theory proposes that inductive category learning occurs via a series of comparisons between current and previous objects. The learner draws on similarities and differences between objects encountered one after another (Carvalho & Goldstone, 2015). The idea is that cognitive systems prioritize which features to encode at a local temporal level and can take in the potential information value of each feature of an object and compare that to the corresponding feature value of the previous object (Carvalho & Goldstone, 2015). This idea is consistent with both behavioral (Jones & Sieck, 2003) and brain imaging data (Schlichting & Preston, 2015; Zeithamova, Schlichting & Preston, 2012; Carvalho, & Goldstone, 2015). Being able to focus one's attention for longer is both important and beneficial because doing so may provide more opportunities to make these types of comparisons in learning situations. It is also well known that training learners with different object sequences and patterns can alter attention and learning outcomes (e.g., Carvalho & Goldstone, 2017).

To summarize, all of these attention and learning findings lead us to put forward the *"non-uniformity-attention hypothesis"* which predicts that babies pay longer attention to object sequences sampled from a non-uniform distribution and therefore have opportunities to build up knowledge about one object, that may then serve to scaffold knowledge about other new objects. Could this be a contribution to the origin of infants' rapid word learning in early life?

Figure 1: A hypothetical set of the first 10 bowls infants encounter.
All three hypothetical infants see the same range of bowls, but in a different order or with a different modal most-frequent bowl.

Putting it all together

If the "*non-uniformity-attention hypothesis*" operates in everyday life and word learning, then practically speaking it cannot depend on the exact order of experiences with certain objects. Sequences matter because infants are constantly making real time comparisons between objects they see, but infants grow up in diverse environments and have different developmental experiences so it is impractical to assume that an infant must see her own blue bowl first, then her family's red bowls, and finally a metal mixing bowl in order to successfully start on the path of learning what a "bowl" is. It is important to recognize that there are individual differences in experience and development. We know that comparison benefits are important to learning, (e.g., Carvalho & Goldstone, 2015) but we hypothesize that repetition and the presence of a single, most frequent object type (in a non-uniform distribution), is what best helps infants make sequential comparisons and learn the generalized idea of something like a "bowl". Figure 1 illustrates this point by comparing three sets of 10 objects. Imagine Infant 1, 2, and 3's first 10 experiences with bowls coincide with the top, middle, and bottom rows, respectively. Infant 1 and Infant 2 see the same modal bowl (small, blue bowl) six times and both infants also see the same range of bowl-types (five unique bowls). While the bowl distribution in rows 1 and 2 are the same, they differ in the order they are presented.

Word learning is a robust process. While not all infants experience the world in the same way or same order, there is evidence that non-uniformity is widespread and infants are overwhelmingly encountering one within-category modal object-type a lot more than other object types (see Table 1). We hypothesize that the features of the modal object type may not matter as much as the fact that there is repetition. For example, while Infant 1 and 2's most frequently appearing bowl is a small, blue bowl, Infant 3's is a medium, red bowl. Infants have a lot to learn from structure in their environments and a robust mechanism is needed to explain what helps infants attend to, listen to, and want to actively sample from their environments. We propose that novel objects sampled from differently structured distributions (uniform vs. non-uniform) and infant encounters with these distributions matter for how infants attend to and learn about new objects and their names.

The present study focuses its measurement on infant attention because attention has obvious and important consequences for learning. The ability to pay attention for longer and actively interact with objects in one's environment not only helps infants see and encounter more, but may positively contribute to learning. For example, toddlers who exhibit high attentional focus generally tend to have better learning outcomes (Dixon, Salley & Clements, 2007). Ideally, we want a mechanism that keeps infants attending to, sampling from, and engaged for an extended time period so that they pay attention to naming instances that coincide with objects as they appear in view (Smith & Yu, 2008). Therefore, to test the "non-uniformity-attention" hypothesis with realistic parameters we will need the following: 1) novel within-category objects for infants to see for the first time 2) a set range of within-category objects and 3) a way to allow infants to experience more objects the longer the distribution holds their attention.

Method

*Participant*s.

80 infants (ages 16-30 months old) will be recruited from the University of Oregon Team Duckling database from the Eugene and Springfield area. This age range is when typically developing toddlers show dramatic improvements in vocabulary and object recognition. Infants will be native English speakers, from a family that speaks predominantly English at home. All participants will complete the MCDI (MacArthur-

Bates Communicative Development Inventories; Fenson et al., 2007) in order to assess their vocabulary development as well as questionnaires about their family, demographics, and general development.

Stimuli.

Objects printed on cards and their names will be novel to infants. Novel objects ensure that no past experience or interest in the objects systematically influence behavior during the study. Images of objects come from the NOUN database (Horst & Hout, 2011) and have been digitally manipulated for color and size permutations. There are two object categories. In a study session, an infant will only encounter one type of object. Code will randomly assign the object category to each infant participant. The reason the experimental design employs two categories is to be able to show that attention measures are not object-related; if there is a significant difference in infant attention between the uniform and non-uniform distribution, that trend should be true for both object categories. All the objects in a category will have the same shape, but vary in size and color (see Figure 2). With three different sizes and six different colors, there will be 18 distinct objects per category. The smallest object will be 3.5 by 2.3 inches, the medium sized object will be 6.5 by 4.3 inches and the largest object will be 8.5 by 5.6 inches. The six colors are: red, blue, green, purple, yellow, and teal. Object shape defines the two categories because the earliest named and learned words tend to be those that are well organized by global shape (Smith et al., 2002). Feature variation on two dimensions (size and color) is used in this study to approximate the multi-dimensional space of the real world. For example, infants' toys often differ in size and color combinations (e.g., different sized and colored building blocks or toy trucks). The objects within a category will be called the same name. “Deef” and “vop” are short, easy to pronounce, and simple to pluralize.

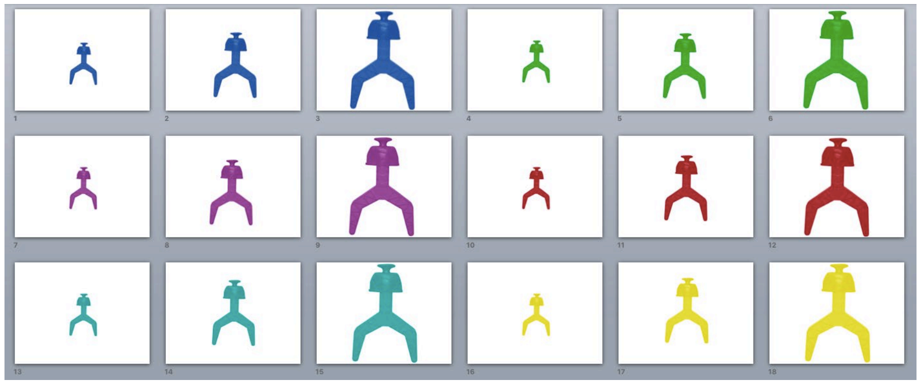

Figure 2: 18 distinct objects varying in color and size.

Design.

Infants will have the opportunity to see up to 40 object trials (four sets of 10 objects). In each distribution and session infants will see the objects one at a time and organized in sets. An infant who makes it through more sets will see a greater variety of objects. For example, the difference between an infant who sees one set and an infant who sees two sets is that the first infant will encounter only five unique object types, while the second infant will encounter ten. The set-organization also ensures that the range of different objects seen at the end of each set of 10 trials, regardless of uniform or non-uniform distribution, is preserved. For example, if two different infants finish seeing all of set 1 in the uniform and non-uniform distributions respectively, they both will have seen an equal number of unique object (see and compare Figures 3 & 4) this ensures that variation did not influence any attention-related differences across sets between the two distributions. It is important that the uniform and non-uniform distribution have the same level of object variety depending on how many sets an infant progresses through so that if we find a significant attention difference we can attribute it distribution type alone and not to the fact that infants saw a greater variety of objects in one distribution. For both distributions, the order of objects within a set will be random and determined by code to model the fact that there is not a universal order in which all infants experience objects and to ensure that object order and any experimenter bias does not influence any measured attention differences. An infant will see object cards presented in either a uniform or non-uniform distribution.

<u>Uniform Distribution:</u> In the uniform distribution condition, an infant will see each object type an equal number of times within a set (10 trials). For example, in the first set, an infant will see each object type (1, 2, 3, 4, 5) two times. In the second set, an infant will see each object type (1, 6, 7, 8, 9) two times and so forth. Note that in the entire 40-object distribution, one object (1) is very weakly "modal" and repeats in each set of 10 trials. All other object cards appear only in one set. This very weakly "modal" object is present in all of the uniform sets as a design choice to preserve the range of unique object types in the uniform and non-uniform distributions. Each infant will have the opportunity to see a maximum of 40 object cards (4 sets of 10 trials).

<u>Uniform distribution:</u>

Set 1: trials 1-10	1	1	2	2	3	3	4	4	5	5
Set 2: trials 11-20	1	1	6	6	7	7	8	8	9	9
Set 3: trials 21-30	1	1	10	10	11	11	12	12	13	13
Set 4: trials 31-40	1	1	14	14	15	15	16	16	17	17

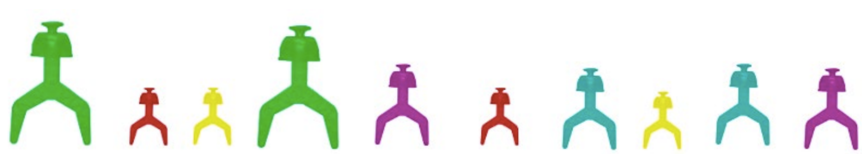

1 2 3 1 4 2 5 3 5 4

Figure 3: Table of four full sets of uniform object trials and a visual example of one set. Numbers indicate unique object types.

<u>Non-uniform Distribution</u>: In the non-uniform distribution condition, an infant will see a modal object (1) six times for every one time she sees other objects (2, 3, 4, 5). Note that within the whole distribution, one object card (1) is "modal" and repeats six times in each set of 10 trials. All other object cards appear only once and in only one set. Each infant will have the opportunity to see a maximum of 40 object cards (4 sets of 10 trials).

<u>Non-uniform distribution:</u>

Set 1: trials 1-10	1	1	1	1	1	1	2	3	4	5
Set 2: trials 11-20	1	1	1	1	1	1	6	7	8	9
Set 3: trials 21-30	1	1	1	1	1	1	10	11	12	13
Set 4: trials 31-40	1	1	1	1	1	1	14	15	16	17

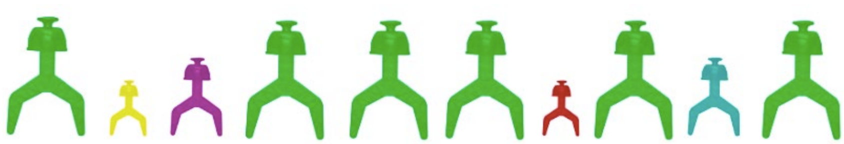

1 2 3 1 1 1 4 1 5 1

Figure 4: Table of four full sets of non-uniform object trials and a visual example of one set. Numbers indicate unique object types.

The dependent measure will be infant attention to objects as measured by 1) the number of objects and 2) the number of full sets an infant views before losing focus or getting off task. Infant looking is widely used and considered an informative behavioral measure of infant attention (Reynolds & Romano, 2016). The *main* dependent measure of infant attention is the number of individual object instances (up to a maximum of 40) that the infant continues to attend to. We are interested in seeing if infants in the non-uniform distribution will get through more objects than in the uniform distribution. However, by considering the overall number of object trials as well as number of object sets (out of four) infants encounter, we will have information on both the full range of objects an infant sees as well as a measure for how far an infant progresses through the four sets. The set organization and measurement allows us to ask questions such as: do all babies make it through all four sets or do babies in one distribution systematically stop before making it through the first set? Our analysis takes into consideration both measures of attention, number of object trials and number of sets, to create a more complete picture of how well each distribution is able to hold infant attention.

Procedure.

In a warm-up, the experimenter and infant will first go through a familiarization period where they will have a free play session for five minutes in order to get familiar with a new person and setting. Next, the experimenter will show the infant a series of pictures of everyday objects and toys (book, ball, hat, etc.) in order to familiarize the infant to the upcoming viewing task. Immediately after this warm-up, the study session will begin.

In the study session, the experimenter will present object cards to infants one at a time. The infant will be seated on parent's lap directly across from the experimenter (see Figure 5 for experiment room setup). After the warm-up period, the experimenter will say "Let's play a game! Ready, set, go!". The experimenter will then pick up the first object card from a concealed area, hold it up to show the infant, name the object with pre-determined script ("Look it's a vop!" or "I have a vop!"), and count to three silently in her head. The experimenter will then put down the object card and grab the next one. The experimenter will repeat this for as many object trials as the infant will pay attention to. A session ends when an infant is no longer "attending" as indicated by looking away from objects and an inability to refocus attention on the viewing task after being prompted verbally by researcher three times.

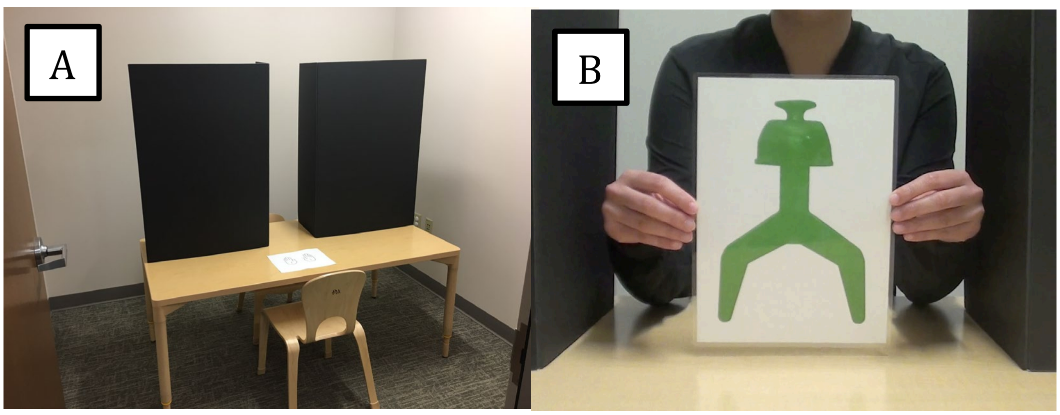

Figure 5: Experimental room setup. (A) Researcher sits at the table on the side closest to wall and the infant sits on the parent's lap directly across from the researcher.
(B) Researcher showing an object card.

For each study session, the in-room experimenter will have the ability to decide when a session ends as indicated by the infant's failure or refusal to begin the next object trial. A session end is indicated by the infant looking away from the objects and a failure to resume the activity after being prompted three times with a pre-script from the experimenter. This script includes the following three sentences: "Focus here!", "Let's keep going!", and "I have more to show you!". After the session, a coder will review the video recording of each session and will code for when the session ended with the same aforementioned guidelines, the number of object trials an infant paid attention to and how many sets of 10 objects the infant attended to. The in-room experimenter's recording of the number of object trials and session time will be compared to the after-session coder's. Any disagreements will be resolved by a third independent review.

Results: Predicted

This study measures infant attention to objects sampled from either a uniform or a non-uniform distribution of object types. According to our "non-uniformity-attention hypothesis" we predict that infants will attend significantly longer to object sequences sampled from a non-uniform distribution. This means that infants will 1) pay attention to a greater mean number of object trials and 2) complete sampling from a greater number of full object sets (i.e., sets of 10 objects) and thereby see a larger range of different unique objects.

Figures 6 and 7 show our main predicted results. Both depict greater attention to objects when sampled from non-uniform distributions. The specific mean value in each experimental condition awaits empirical demonstration, but we predict that the values in the non-uniform condition will be closer to the maxima (40 objects, 4 sets) and the values in the uniform condition will be closer to the minima (0 objects, <1 set).

Figure 6 shows the predicted results for the *mean number of object trials infants attend to.* We predict that infants will pay attention to more objects in the non-uniform distribution and fewer objects in the uniform distribution. Figure 7 shows the predicted results for the *mean number of full object sets infants attend to.* We predict that infants will pay attention to more sets of objects in the non-uniform distribution and perhaps not even finish a full set of objects in the uniform distribution.

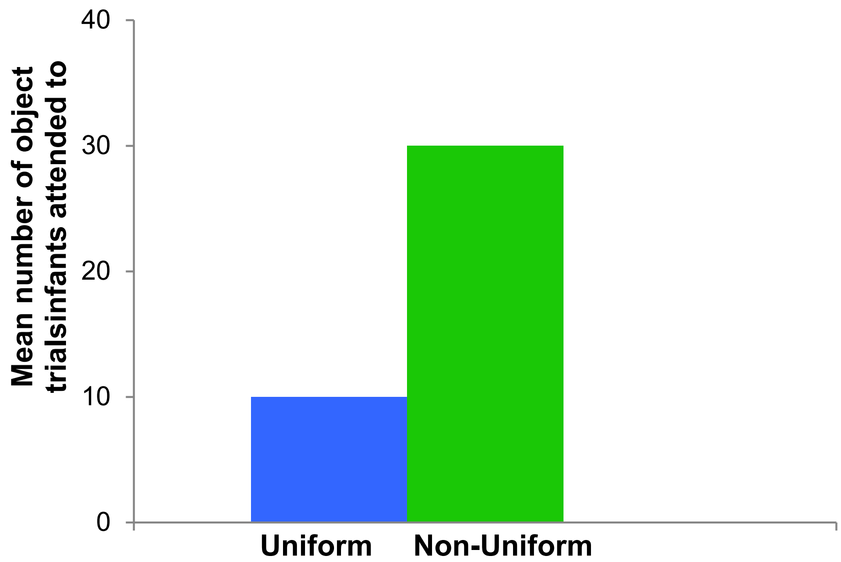

Figure 6: Predicted mean number of object trials infants attended to.

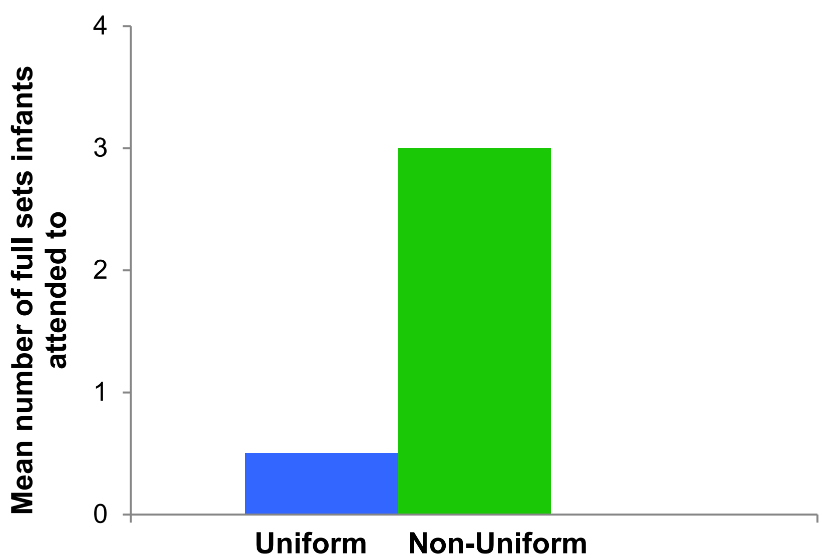

Figure 7: Predicted mean number of full sets infants attended to.

Discussion

If it is true that the structure of the non-uniform object distribution is better able to hold infant attention, then infants will pay attention to more object trials and should get through more full object sets than with the standard uniform distribution. Non-uniform distributions are common in the real-world and in infants' everyday environments, and infants are readily and robustly able to learn in these environments (Smith et al., 2018). Therefore it is possible that infants may be prepared with internal learning mechanisms well-matched to these non-uniform types of data which may help keep them attending for longer periods of time. The non-uniform distribution may provide a helpful balance in its structure, which is characterized by both repetitions of a modal object as well as variability in the range of object examples. A recurring modal object may serve as an anchor that infants can refer to as they encounter more new objects, thus instantiating the idea that perhaps "something old helps infants pay attention to something new". Seeing a larger range of objects (that differ in color and size but share a shape) helps create an understanding that shape is what characterizes an object category (see Smith et al., 2002; Perry et al., 2010). Objects in our world are organized into categories based on the feature of global shape. Attending to shape while others features like color and size change is important to learn the names of objects. The sequential attention theory proposes that a mechanism involving comparing and contrasting last seen objects with newly seen objects is a way of learning. Consequently, if an infant can pay attention for longer, they can create more learning opportunities for themselves both in moment-to-moment instances as well as cumulatively over time because they see a larger range of objects. Being able to pay attention to objects sampled from distributions that mirror real-life, repetitive, and non-uniform statistical structure may have helpful learning consequences.

How does this influence how we study infant learning in the lab? If infants do pay significantly longer attention to the non-uniform distribution than the uniform distribution, this has important methodological consequences for future studies about early infant learning abilities. The vast majority of previous in-lab statistical word learning and developmental studies use uniform distributions. One possibility is that results from studies that present young learners with instances in uniform distributions underestimates their ability to attend and learn. This raises the question, if infants in these previous studies with only uniform distributions had paid more attention, would they have learned more? This issue is especially relevant to infant studies, given the field's high dropout rates (Bergmann et al., 2018). "Dropout rates" refer to the proportion of infants who enter a study, but do not complete it. Reasons for dropping out include a failure to engage in the study task due to a crying fit or temper tantrum. Infants can be temperamental and collecting data from them is time-consuming. Relatively high dropout rates are fairly common for in-lab developmental psychology studies and it is worth considering why this is the case. It is possible that existing infant learning experimental paradigms that at first appear to "fail", i.e. appear to be ill-suited to measure abilities of interest to developmental researchers, boil down to a failure in ability to capture fragile infant attention. Could presenting stimuli in non-uniform distributions help with this challenge? If so, this would be an exciting and productive advance for developmental studies aiming to make discoveries about how young infants learn.

This discussion raises questions about how much we actually know about structure in the real-world. While recent research advances have increased our understanding of statistical structure in infant egocentric views, there is still much to explore in structures present in the environments of young children. How does this structure change with age? Does it vary by category, for example is it more skewed for some kinds of objects than others? How much do individual experiential differences contribute to differences in attention? This leads us to an important point on the limits of generalization that exist in the proposed research study. The present study's design does not take into consideration individual infants' differences in experience, socio-economic status, and home-lives. We recognize that these are important factors to consider for experimental design purposes. Our current study seeks to compare one non-uniform distribution to one uniform distribution. The lack of previous research regarding this specific developmental problem forces us to begin on a basic level and ask the essential question of whether non-uniformity can better capture infant attention than the standard lab use of uniform distributions. We recognize that there is much room for expansion on this original research question and to take into consideration individual infants differences. Perhaps the problem is not even one of modeling in-lab distributions to match real-world distributions, but one of modeling in-lab distributions to specific and different infants' experiences.

In summary, results from the present study could indicate a possibility that we need to rethink how we design in-lab studies that aim to characterize infants' statistical learning abilities. What can we aim to learn about infant attention and learning development from studies that use statistical distributions so markedly different from what infants are actually experiencing in the real-world? The present study can not only provide important insight into what types of distributions best capture infant attention, but also has the potential to motivate changes in developmental study methodology.

Understanding the interactions among attention, the sequence in which information is studied, and what information is presented, can ultimately lead to a better appreciation of how learning unfolds in real time (Carvalho & Goldstone, 2017). The present study has the potential to contribute to understanding how infants learn the names of new objects. Previous research indicates that the names of objects that infants learn first are those that are highly prevalent in their view and everyday lives (Clerkin et al., 2016). As infants get older, physical developments increasingly allow for more agency and an ability to move around and manipulate the physical surroundings. Consequently, this increase in mobility causes an increase in visual exposure to not only a greater number of objects, but also a greater number of less familiar objects. The present study was designed to find out what type of object distribution may systematically and significantly better hold infant attention. For practical reasons, the present study has infants sitting still while sampling from the object distributions, but we recognize that in the real-world developing infants spend increasing amounts of time moving around in and actively sampling from their environments. This is an important developmental problem because how infants move through their physical spaces in concordance with the visual processes at work dictate that infants will frequently encounter objects in informal and natural sequences.

The current study is a work-in-progress. In the context of this study, more work can be done to further explore infant attention and to tie attention to memory and learning. The immediate next step is to complete data collection for the present study.

Next, we plan to include memory and learning tests after the object presentation period. Are there concrete learning effects between sampling novel objects from a uniform vs. non-uniform distribution? We predict that if infants pay longer attention to novel objects sampled from a non-uniform distribution they could learn the names of the novel objects better than from a uniform distribution. Repetition and variability are key to learning new words and by paying attention to a distribution for longer, an infant increases her chances of encountering repeat objects as well as a greater variety of object types. We also predict that the non-uniform distribution could help infants learn to generalize and sort new in-category objects with greater accuracy. Seeing multiple exemplars repeatedly may help an infant learn the characteristics of an object category better than if they just encountered a few examples. A solidified understanding of the characteristics of what defines an object category may then help an infant better identify new in-category objects.

Several theoretical issues involving attention, memory, and learning in developmental psychology are relevant to the present study. The study of rates of infant attention to novel objects presented in different distributions (uniform vs. non-uniform) has real-world implications. With enough data collected from this study, detailed attributes of object distribution(s) that best capture infant attention may be clearly delineated. The goal is to find if an “attentional sweet-spot” exists, perhaps with a certain number of objects presented in a sequence involving a specific range or rate of modal object repetition. The information derived from this study could also have profound effects on both formal and informal teaching methods. If we are to develop a predictive and generalizable model of object distribution that best captures infant visual attention, it is helpful to take a step back and put this project into perspective of structure in infant environments.

This research matters not only for improving our understanding of infant attention and how different object sequences affect attention, but also has potentially powerful methodological influence. This study takes on the challenge of how to examine the consequences of recently discovered statistical patterns in infants' everyday lives head on and brings these non-uniform patterns into the lab. The study method moves beyond standard in-lab study approaches and tests for a hypothesized sweetspot in non-uniform distributions between repetition and novelty. This sweetspot represents an exciting possible solution to reduce the field's high dropout and fuss out rates in infant behavioral studies. The ultimate question is what role does attention play in learning? The present study will begin to answer this question and will yield new insight into the effects for understanding infants' attention and statistical learning abilities in the context of early word acquisition. In conclusion, creating a better understanding of object distributions that best captivate infant visual attention is an important and fundamental step to understanding cognitive development at a broader level.

Works Cited

Bergmann, C., Tsuji, S., Piccinini, P. E., Lewis, M. L., Braginsky, M. B., Frank, M. C., & Cristia, A. (2018). Promoting replicability in developmental research through meta-analyses: Insights from language acquisition research. Child Development. Advance online publication. doi:10.1111/cdev.13079.

Carvalho, P. F., & Goldstone, R. L. (2015). The benefits of interleaved and blocked study: Different tasks benefit from different schedules of study. *Psychonomic Bulletin & Review,22*(1), 281-288. doi:10.3758/s13423-014-0676-4

Carvalho, P. F., & Goldstone, R. L. (2017). The sequence of study changes what information is attended to, encoded, and remembered during category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition,43*(11), 1699-1719. doi:10.1037/xlm0000406

Castelhano, M. S., Mack, M. L., & Henderson, J. M. (2009). Viewing task influences eye movement control during active scene perception. *Journal of Vision,9*(3), 6-6. doi:10.1167/9.3.6

Clerkin, E. M., Hart, E., Rehg, J. M., Yu, C., & Smith, L. B. (2016). Real-world visual statistics and infants first-learned object names. *Philosophical Transactions of the Royal Society B: Biological Sciences,372*(1711), 20160055. doi:10.1098/rstb.2016.0055

Dixon, W. E., Salley, B. J., & Clements, A. D. (2006). Temperament, distraction, and learning in toddlerhood. *Infant Behavior and Development,29*(3), 342-357. doi:10.1016/j.infbeh.2006.01.002

Fausey, C. M., Jayaraman, S., & Smith, L. B. (2016). From faces to hands: Changing visual input in the first two years. *Cognition,152*, 101-107. doi:10.1016/j.cognition.2016.03.005

Fennell, C. T. (2011). Object Familiarity Enhances Infants' Use of Phonetic Detail in Novel Words. *Infancy,17*(3), 339-353. doi:10.1111/j.1532-7078.2011.00080.x

Fenson, L., Marchman, V. A., Thal, D. J., Dale, P. S., Reznick, J. S., & Bates, E. (2007). MacArthur-Bates Communicative Development Inventories: User's guide and technical manual (2nd ed.). Baltimore, MD: Brookes.

Fisher, A.V., & Kloos, H. (2016). Development of selective sustained attention: The role of Executive Functions. In J. A. Griffin, P. McCardle, & L. Freund (Eds.), Executive Function in Preschool-age Children: Integrating Measurement, Neurodevelopment, and Translational Research, 215-237. Washington, DC, US: American Psychological Association.

Fiser, J., & Aslin, R. N. (2002). Nonlinear partial differential equations and applications: From the Cover: Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences,99*(24), 15822-15826. doi:10.1073/pnas.232472899

Fisher, A. V., Goodwin, K. E., & Seltman, H. (2014). Visual Environment, Attention Allocation, and Learning in Young Children: When too much of a good thing may be bad. *Psychological Science,25*(7), 1362-1369.

Gerken, L., Balcomb, F. K., & Minton, J. L. (2011). Infants avoid 'labouring in vain' by attending more to learnable than unlearnable linguistic patterns. *Developmental Science,14*(5), 972-979. doi:10.1111/j.1467-7687.2011.01046.x

Horst, J. S., Parsons, K. L., & Bryan, N. M. (2011). Get the Story Straight: Contextual Repetition Promotes Word Learning from Storybooks. *Frontiers in Psychology,2*. doi:10.3389/fpsyg.2011.00017

Horst, J. S. & Hout, M. C. (in press). The Novel Object and Unusual Name (NOUN) Database: a collection of novel images for use in experimental research. *Behavior Research Methods.* Manuscript accepted for publication.

Hunter, M. A., & Ames, E. W. (1988). A multifactor model of infant preferences for novel and familiar stimuli. In C. Rovee-Collier & L. P. Lipsitt (Eds.), *Advances in infancy research,* Vol. 5, pp. 69-95). Westport, CT, US: Ablex Publishing.

Jayaraman, S., Fausey, C. M., & Smith, L. B. (2015). The Faces in Infant-Perspective Scenes Change over the First Year of Life. *Plos One,10*(5). doi:10.1371/journal.pone.0123780

Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G. F., Mcclure, S. M., Wang, J. T., & Camerer, C. F. (2008). The Wick in the Candle of Learning: Epistemic Curiosity Activates Reward Circuitry and Enhances Memory. *SSRN Electronic Journal*. doi:10.2139/ssrn.1308286

Kidd, C., & Hayden, B. (2015). The Psychology and Neuroscience of Curiosity. *Neuron,88*(3), 449-460. doi:10.1016/j.neuron.2015.09.010

Lee, D. K., Cole, W. G., Golenia, L., & Adolph, K. E. (2017). The cost of simplifying complex developmental phenomena: A new perspective on learning to walk. *Developmental Science*. doi:10.1111/desc.12615

Malcolm, G. L., Rattinger, M., & Shomstein, S. (2016). Intrusive effects of semantic information on visual selective attention. *Attention, Perception, & Psychophysics,78*(7), 2066-2078. doi:10.3758/s13414-016-1156-x

Mather, E., & Plunkett, K. (2009). Learning Words Over Time: The Role of Stimulus Repetition in Mutual Exclusivity. *Infancy,14*(1), 60-76. doi:10.1080/15250000802569702

Mcmurray, B. (2007). Defusing the Childhood Vocabulary Explosion. *Science,317*(5838), 631-631. doi:10.1126/science.1144073

Perry, L. K., Samuelson, L. K., Malloy, L. M., & Schiffer, R. N. (2010). Learn Locally, Think Globally. *Psychological Science,21*(12), 1894-1902. doi:10.1177/0956797610389189

Piantadosi, S. T. (2014). Zipf's word frequency law in natural language: A critical review and future directions. *Psychonomic Bulletin & Review,21*(5), 1112-1130. doi:10.3758/s13423-014-0585-6

Piantadosi, S. T., Kidd, C., & Aslin, R. (2014). Rich analysis and rational models: Inferring individual behavior from infant looking data. *Developmental Science,17*(3), 321-337. doi:10.1111/desc.12083

Reynolds, G. D., & Romano, A. C. (2016). The Development of Attention Systems and Working Memory in Infancy. *Frontiers in Systems Neuroscience,10*. doi:10.3389/fnsys.2016.00015

Roder, B. J., Bushneil, E. W., & Sasseville, A. M. (2000). Infants Preferences for Familiarity and Novelty During the Course of Visual Processing. *Infancy,1*(4), 491-507. doi:10.1207/s15327078in0104_9

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition,70*(1), 27-52. doi:10.1016/s0010-0277(98)00075-4

Saffran, J. R., & Wilson, D. P. (2003). From Syllables to Syntax: Multilevel Statistical Learning by 12-Month-Old Infants. *Infancy,4*(2), 273-284. doi:10.1207/s15327078in0402_07

Salakhutdinov, R., Torralba, A., & Tenenbaum, J. (2011). Learning to share visual appearance for multiclass object detection. *Cvpr 2011*. doi:10.1109/cvpr.2011.5995720

Schwab, J. F., & Lew-Williams, C. (2016). Repetition across successive sentences facilitates young children's word learning. *Developmental Psychology,52*(6), 879-886. doi:10.1037/dev0000125

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. doi:10.1111/1467-9280.00403

Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition,106*(3), 1558-1568. doi:10.1016/j.cognition.2007.06.010

Smith, L. B., Jayaraman, S., Clerkin, E., & Yu, C. (2018). The Developing Infant Creates a Curriculum for Statistical Learning. *Trends in Cognitive Sciences,22*(4), 325-336. doi:10.1016/j.tics.2018.02.004

Spain, M., & Perona, P. (2002). Measuring and Predicting Importance of Objects in Our Visual World. *Computation and Neural Systems,*1-8.

Twomey, K. E., Ranson, S. L., & Horst, J. S. (2013). Thats More Like It: Multiple Exemplars Facilitate Word Learning. *Infant and Child Development,23*(2), 105-122. doi:10.1002/icd.1824

Wonnacott, E., Brown, H., & Nation, K. (2017). Skewing the evidence: The effect of input structure on child and adult learning of lexically based patterns in an artificial language. *Journal of Memory and Language,95*, 36-48. doi:10.1016/j.jml.2017.01.005