Article

FIRST LANGUAGE

First Language
1–30

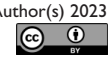

DOI: 10.1177/01427237221150070
journals.sagepub.com/home/fla


How the characteristics of words in child-directed speech differ from adult-directed speech to influence children's productive vocabularies

Gary Jones
Nottingham Trent University, UK

Francesco Cabiddu
Cardiff University, UK

Doug J. K. Barrett
University of Leicester, UK

Antonio Castro
Bethany Lee
Nottingham Trent University, UK

Abstract

Child-directed speech has long been known to influence children's vocabulary learning. However, while we know that caregiver utterances differ from those directed at adults in various ways, little is known about any differences in the lexical properties of child-directed and adult-directed utterances. We compare over half a million word tokens from adult speech directed at children (from caregiver–child transcriptions) to the same quantity directed at adults. We show that child-directed speech contains greater numbers of words that are lower in phonemic length, higher in frequency, lower in phonotactic probability, and higher in neighborhood density than adult-directed speech; furthermore, child-directed speech explains over twice the variability of children's productive noun vocabularies than adult-directed speech. These findings indicate that children's word production is clearly influenced by the characteristics of the words

Corresponding author:
Gary Jones, Department of Psychology, Nottingham Trent University, 50 Shakespeare Street, Nottingham NG1 4FQ, UK.
Email: gary.jones@ntu.ac.uk

spoken directly to them and that researchers need to be wary of using adult-directed language corpora when calculating lexical measures.

Keywords
Child-directed speech, word learning, motherese, caregiver speech, vocabulary

Introduction

Child-directed speech differs from adult-directed speech in ways that influence children's language acquisition, such as exaggerated intonation and shorter utterances (e.g. Saxton, 2009). The lexical characteristics of child-directed speech also influence child language, with early word production being influenced by factors such as spoken frequency of occurrence and word length (e.g. Jones et al., 2021). However, it is unclear whether the lexical characteristics of words learned in early childhood are more similar to the lexical characteristics of words in child-directed speech than adult-directed speech, and consequently, whether such lexical characteristics facilitate early word learning more so for child-directed speech than adult-directed speech. In this article, we provide a detailed comparison of the characteristics of words used in the vocabularies of child-directed and adult-directed speech, together with how those characteristics map onto children's productive vocabularies. We will show how children's early expressive vocabularies are influenced by the lexical characteristics of child-directed speech in particular, suggesting that the characteristics of words spoken to children promote their production of some words over others.

Adults change the way they talk depending on who they are speaking to. For example, when native-speaking adults talk to other adults for whom the native language is foreign, spoken language is simplified (e.g. Rodriguez-Cuadrado et al., 2018). It should therefore come as no surprise that if an adult is speaking to a child, their spoken language will be different to that used when speaking to another adult. This is particularly the case when addressing infants (e.g. Saxton, 2009; Suttora et al., 2017). First, intonation is more pronounced such that communicative intent is clearer (e.g. Fernald, 1989), with child-directed speech being characterized by a slower speech rate and greater variability in pitch (e.g. Thiessen et al., 2005). These alterations to typical adult speech are believed to aid the discriminability of vowels and potentially consonants (Liu et al., 2003) with a reduction in speech rate potentially being a key factor (McMurray et al., 2013). Second, child-directed speech also involves fewer words than adult-directed speech (Saxton, 2009), with speech length being adapted to suit the linguistic competence of the child (Roy et al., 2009). Third, sentence structure and word content are changed from what would be used when talking to other adults. For example, Furrow et al. (1979) showed that speech directed toward 18- to 27-month-old children had greater numbers of *wh*-questions and fewer declaratives than adult-directed speech, while You et al. (2021) showed that child-directed speech allowed causal meaning to be inferred based on the co-occurrence of neighboring words, which was not possible with adult-directed speech.

The changes in the spoken language used when adults address children rather than other adults may well influence children's vocabulary acquisition. Weisleder and Fernald (2013) measured both child-directed speech and adult-directed (overheard) speech involving 19-month-old infants within 29 families, finding only the former to be linked to infant vocabulary size 6 months later. Changes in the length of spoken utterances seems to have an effect on child productive vocabularies, with shorter utterances directed to 18-month-old infants relating to larger productive vocabularies at 27 months (Furrow et al., 1979), while for speech directed to older infants (18–29 months old) the relationship reverses (Hoff & Naigles, 2002).[1] With regard to intonation, 7-month-old infants are more able to identify word boundaries in a continuous stream of nonsense words when the words are spoken with exaggerated prosody that is consistent with child-directed speech, rather than that of adult-directed speech (Thiessen et al., 2005). Since identifying words in a continuous stream of speech is part of the challenge for word learning, one may expect the exaggerated prosody of child-directed speech to facilitate vocabulary acquisition. There are also differences in word use across child-directed and adult-directed speech, and these differences may also have consequences for children's word learning. For example, Hills (2013) measured contextual diversity and word repetitions in child-directed speech and adult-directed speech, finding that child-directed speech was less diverse and more repetitive than adult-directed speech. Moreover, children's word learning was better predicted by child-directed speech.

Current study

While there are many examples of how child-directed speech may benefit children's vocabularies, few if any compare the lexical properties of child-directed and adult-directed speech. The current study examines four well-established lexical characteristics (word length, frequency of occurrence, neighborhood density, and phonotactic probability) to see if they differ across child-directed and adult-directed speech and to what extent any differences are predictive of children's early vocabularies.

A number of metrics have been used to characterize the lexical properties of the early words in child vocabularies: word length,[2] word frequency, neighborhood density, and phonotactic probability. Not surprisingly, early vocabularies contain more short than long words (e.g. Maekawa & Storkel, 2006; Storkel, 2009). This is consistent with a wealth of memory and related literature where both children and adults have greater difficulty recalling long words and long nonsense words (e.g. Baddeley et al., 1975; Gathercole & Baddeley, 1989). For example, numerous studies show that children are more accurate at repeating nonsense words when the nonsense words are short rather than long (e.g. Jones et al., 2010; Jones & Witherstone, 2011). However, we are unaware of whether child-directed speech contains a proportionally larger number of short than long words relative to adult-directed speech. It seems plausible that this would be the case on the basis of mean length of utterance being smaller in child-directed speech than adult-directed speech (e.g. Bernstein Ratner & Rooney, 2001), but as far as we are aware, no study has compared word length between the two types of speech.

Child vocabularies are also dominated by words that occur frequently in the native language. For example, Huttenlocher et al. (1991) showed that as the frequency by which

particular content words occurred increased in parental speech, the age at which their children produced those same words decreased. A similar pattern is found for verbs (Naigles & Hoff-Ginsberg, 1998) and more generally across word categories. Goodman et al. (2008) examined common nouns, people words, verbs, adjectives, closed-class, and other words using 3.8 million tokens from mother–child transcripts on CHILDES (MacWhinney, 2000), showing that higher parental frequency of words in each category led to earlier child production of those words. Similarly, Swingley and Humphrey (2018) showed that words occurring with a high rather than low frequency in the maternal input were far more likely to be present in the child's productive vocabulary. In the work above, word frequency is calculated on the basis of child-directed speech. In other work, word frequency also influences vocabulary acquisition when it is measured based on adult-directed written language (e.g. Stokes, 2010). That said, it is not clear whether child-directed speech contains greater numbers of high frequency words over adult-directed language. As far as we are aware, the only work to compare the two was Hayes and Ahrens (1988) who found no word frequency differences between speech directed at young children (up to 2 years of age) and speech aimed at older children (2–12 years). However, their analysis used a relatively small number of speech tokens and needs to be verified using larger samples.

Child vocabularies consist of proportionally more dense neighborhood words than adult vocabularies (Coady & Aslin, 2003). For any given word, a neighbor is another word that differs by the addition, deletion or substitution of one sound, for example, *slit, it*, and *lot* are neighbors of *lit*. Based on examination of MacArthur Child Development Inventories (MCDIs, checklists that allow caregivers to indicate words known and words produced by infants), Storkel (2009) showed that children aged 16–28 months knew a greater number of nouns with many neighbors than they did those with few neighbors. The effect of neighborhood density appears to be strong. For example, Stokes (2010) examined all monosyllabic words in the MCDI and found 47% of the variance in 24- to 30-month-old children's vocabularies was accounted for by their neighborhood density in adult language. Knowledge of similar sounding words scaffold the learning of a novel word (see Jones et al., 2021); if caregiver speech contains a greater proportion of dense neighborhood words then children are receiving linguistic input that is likely to aid their word learning.

A limitation of prior work examining the influence of neighborhood density on child-directed speech is that previous work has computed neighborhood density using either adult-directed written language or child-directed speech. When computing neighborhood density using adult-directed written language (e.g. from a 20,000 word adult dictionary; Storkel, 2009), it is a significant predictor of child vocabularies (Stokes, 2010; Storkel, 2009). When calculated using child-directed speech (using child-directed speech from Brent & Siskind, 2001), it no longer appears to be (Swingley & Humphrey, 2018), although a network analysis of children's early productive vocabularies showed that the number of neighbors (calculated from caregiver utterances) was positively related to when a word enters the child's productive lexicon (Carlson et al., 2014). One potential confound across many of these studies is the use of child-directed *speech* versus adult-directed estimates based on *written* text, the latter being a clear departure from the language that children actually hear. In the current study we will compare the effect of

neighborhood density on child vocabularies when computing densities based on both child-directed and adult-directed speech.

Phonotactic probability has also been shown to hold an influence over children's word learning. Words that are high in phonotactic probability are those that have common sound sequences (traditionally measured at the biphone level) within them; whereas words that are low in phonotactic probability have rarer sound sequences. Storkel (2009) used MCDIs to show that the receptive vocabularies of young children chiefly consisted of words that were low in phonotactic probability. Studies using quick incidental learning also show that young children are more likely to identify a novel object when it is labeled with a novel low phonotactic probability word than when paired with a novel high phonotactic probability word (e.g. Storkel & Lee, 2011). That said, once again there is an issue in measure calculation: the cited work used adult-directed written language corpora for calculations. When child-directed speech was used, phonotactic probability was not a significant predictor of early word learning when word frequency was accounted for (Swingley & Humphrey, 2018). We will therefore apply the same corpora for neighborhood density to also examine phonotactic probability.

In summary, children's early vocabularies contain large numbers of short words, high frequency words, and low phonotactic probability words, and a greater proportion of high neighborhood density words, but it is not clear whether this is due to being exposed to more similar distributions in child-directed speech relative to adult-directed speech.

Note that an alternative view could be that the linguistic input has a secondary influence on children's vocabulary learning. Under this view, the child's own cognitive apparatus (e.g. phonological short-term memory, Gathercole & Baddeley, 1989) is the primary constraint on the type of words that can be learned. This suggests that early word learning would be relatively uniform, such as children learning short words regardless of the length of words addressed to them. However, as we will see, children's early word learning shows similar distributions to the linguistic input they hear (see, for example, Jones et al., 2021), suggesting that the linguistic input plays a large role in vocabulary acquisition.

Our primary focus, therefore, is to examine whether children's learning of words is influenced by the characteristics of the words that they hear. That is, does child-directed speech predominantly contain short words, high frequency words, low phonotactic probability words, and (proportionally) dense neighborhood words; and how predictive are these measures of children's productive vocabularies? Our secondary focus is to also examine adult-directed speech. If, as we argue, child vocabularies are in part driven by the lexical characteristics of words in child-directed speech, then we should expect those characteristics to be more predictive of child vocabularies when they are measured within child-directed speech than when they are measured in adult-directed speech. Examining adult-directed speech will therefore enable us to determine whether word characteristics of adult-directed speech are different to that of child-directed speech and have different predictive powers in relation to child vocabulary learning, since this could add weight to any evidence that children's word learning is influenced by the lexical characteristics of the speech they hear. In addition, because previous work has shown conflicting results based on the corpora used to calculate word frequency, phonotactic probability, and neighborhood density, computing these using both child-directed and adult-directed speech will show how these differ on the basis of the measurement corpora.

The remainder of this article is as follows. First, we outline our methodology and the child-directed and adult-directed speech corpora from which our four measures (word length, word frequency, neighborhood density and phonotactic probability) will be derived together with the corpora from which child vocabularies will be derived. We then detail comparisons across the child-directed and adult-directed speech corpora for all four measures together with the utility of each corpus in predicting child productive vocabularies (productive rather than receptive vocabulary is examined because we use child speech corpora rather than MCDIs). Finally, we discuss our results in the context of current literature.

Method

The speech corpora

Child-directed speech is aggregated from six CHILDES British English corpora, extracting only the child-directed utterances when the children were 2–3 years of age: Belfast (Henry, 1995), Thomas (Lieven et al., 2009), Tommerdahl and Kilpatrick (2013), Wells (1981), Forrester (2002), and Lara (Rowland & Fletcher, 2006).

Adult-directed speech is taken from the British National Corpus (BNC, 2007). The BNC is a monolingual British English corpus containing approximately 100 million orthographic words of spoken (10%) and written language (90%). The spoken BNC was used in the present study. This consists of transcriptions of unscripted informal conversations balanced for age, region, context, and social class and also formal conversations, for example, from meetings.

Child productive vocabularies were taken from the child utterances of the Manchester corpus (Theakston et al., 2001), which involve 12 sets of mother–child interactions spanning 1 year between (average) child ages 22–34 months. Recordings for each mother–child pairing were taken twice every 3 weeks and orthographically transcribed. We chose not to use the maternal utterances from these corpora as our child-directed speech to eliminate any overlap in vocabulary use that may influence our analyses (i.e. the context of mother–child conversations may involve similar vocabulary).

For all corpora, compound nouns were separated into their constituent words (e.g. *merry-go-round* as three words: *merry, go*, and *round*). The rationale for this was (1) our corpus analysis involves 2- to 3-year-old children who are able to both produce novel compounds and are able to determine that words in compounds have different roles (e.g. interpreting cheese + knife as a knife to cut cheese rather than a knife made out of cheese), suggesting children of this age understand the individual words in compounds (Clark et al., 1985); (2) our corpus analyses rely on transcriptions of compounds that vary within and across transcribers (e.g. Christmas tree and Christmas + tree are transcribed), so separating out compounds make our analyses more consistent; and (3) many compounds, such as *merry-go-round*, would have zero neighbors if treated as one entity. To show that the way in which we treat compounds does not influence our results, in Appendix 1 we duplicate our analyses when excluding all compounds, with the results being identical to those in the main article.

Given that the corpora involve orthographic transcriptions and three of our four measures require phonological transcriptions, we converted all orthographic forms to

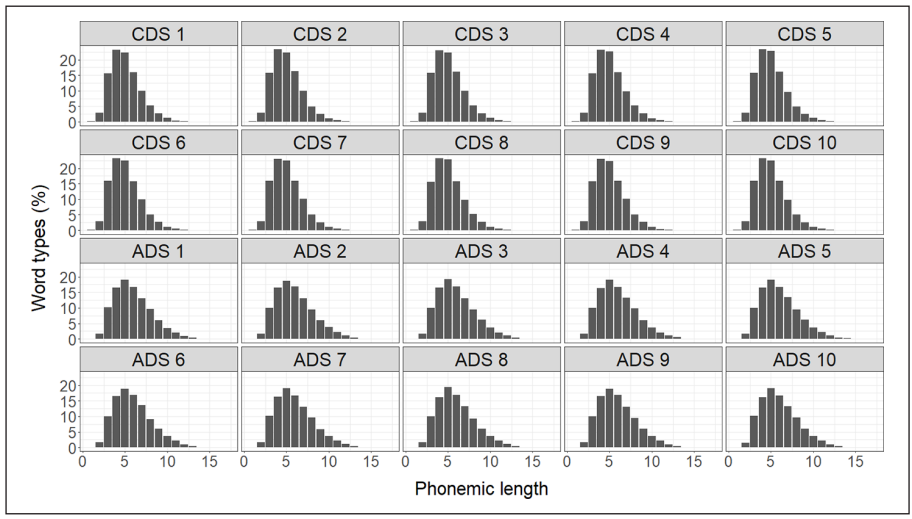

Figure 1. Distribution of Word Lengths across 10 Random Samples from the Child-Directed Speech (Upper Two Panels, Labeled CDS 1-10) and Adult-Directed Speech (Lower Two Panels, Labeled ADS 1-10) to Illustrate that the Distributions from Each Sample Are Almost Identical.

phonological forms using the CMU Pronouncing Dictionary (http://www.speech.cs.cmu.edu/cgi-bin/cmudict) containing over 134,000 American English phonetic transcriptions of words. Commonly occurring words in the children's speech corpora that were not present in the dictionary were manually added together with their phonemic equivalent, as were words that were transcribed differently to how they were spoken (e.g. won't in place of willn't). Orthographic utterances were ignored when they contained a word that did not have a corresponding entry in the dictionary. Unless otherwise stated, all quantities reported for corpora are based on the phonological form. Note that by considering the phonological form only, some orthographically different word types (e.g. *to, two, too*) are treated as one word type.

The complete set of children's phonemically-transcribed utterances comprised 531,940 word tokens, whereas the child-directed speech comprised 1,431,271 word tokens and the adult-directed speech comprised 6,974,576 word tokens. Differences in sample size are problematic for comparisons across corpora because distributional properties change as sample size increases – in particular, increases in number of word types decline as sample size increases (see Montag et al., 2018). We therefore created 10 random samples of the child-directed speech and adult-directed speech to each match the children's utterances for number of word tokens. This was accomplished by randomly selecting utterances from the child-directed speech and adult-directed speech corpora until the number of word tokens matched that of the children. The distributions of word length, word frequency, phonotactic probability and neighborhood density were very similar across all child-directed speech samples and across all adult-directed speech samples (see Figure 1 for word length distributions across different random samples) and we therefore used one random sample from each for subsequent reporting and analyses.

Vocabulary examination

Data manipulation and analysis were carried out using the R programming language (version 3.5.2; R Core Team, 2018).

All corpora included tags to indicate noun plural and so for all word-type analyses, noun plurals were reduced to their singular form (e.g. *car* for *cars*).

Calculations were produced for each of our four measures. Phonemic word length does not change based on the corpora used and was the phonemic length of each word type. The remaining three measures were computed for both child-directed speech and adult-directed speech.

Word frequency was the number of occurrences of each word type within the respective corpora, computed as log10 frequencies.

Neighborhood density was defined as stress-unmarked phonological neighborhood density, referring to the unweighted count of words that differed from a given word by one phoneme (i.e. neighborhood density was a count of all word types that were neighbors of the target word in the corpora, in line with the majority of studies involving neighborhood density).

Phonotactic probability was defined as stress-unmarked word-average biphone probability, referring to the weighted likelihood of the occurrence of ordered phoneme pairs that were present in a given word (i.e. phonotactic probability reflected how often the biphone occurs in the corpora, in line with the majority of studies involving phonotactic probability). This was computed by the second author using the formulae from the Irvine Phonotactic Online Dictionary (version 2.0; IPhOD, Vaden et al., 2009, see http://www.iphod.com/) obtained from the main author of the dictionary (Vaden, personal correspondence, 21/07/2019).

All words, together with their lexical measures as computed from the child-directed corpora and as computed from the spoken BNC can be downloaded here: https://www.ntu.ac.uk/research/groups-and-centres/groups/language-literacy-and-psycholinguistics.

Results

All analysis scripts can be found here: https://doi.org/10.17605/OSF.IO/TCQU5. We first outline the descriptive statistics for each corpus and then provide analyses for phonemic word length, word frequency, neighborhood density, and phonotactic probability. For our analyses, three of the four aforementioned measures (word frequency, phonotactic probability, and neighborhood density) are computed using child-directed speech and adult-directed speech, since these may vary depending on the particular corpora used to compute them. Note that in these analyses, all word categories are considered. Please see the Appendix 2 for plots that are broken down by word category (noun, verb, adjective, and 'other'), which show that there is little difference when considering all word categories as a whole versus when breaking them down by category.

Descriptive statistics for corpora

Table 1 shows descriptive statistics for child productive vocabularies, child-directed speech, and adult-directed speech. There are two notable aspects: (1) by examining

Table 1. Descriptive statistics for phonological word types/tokens for each corpus (child productive vocabularies, child-directed speech, adult-directed speech).

Corpora	Total word types	Total word tokens	Sample word types	Sample word tokens
Child vocabularies	4051	531,940	NA	NA
Child-directed speech	8573	1,431,271	6503	531,940
Adult-directed speech	29,539	6,974,576	13,942	531,940

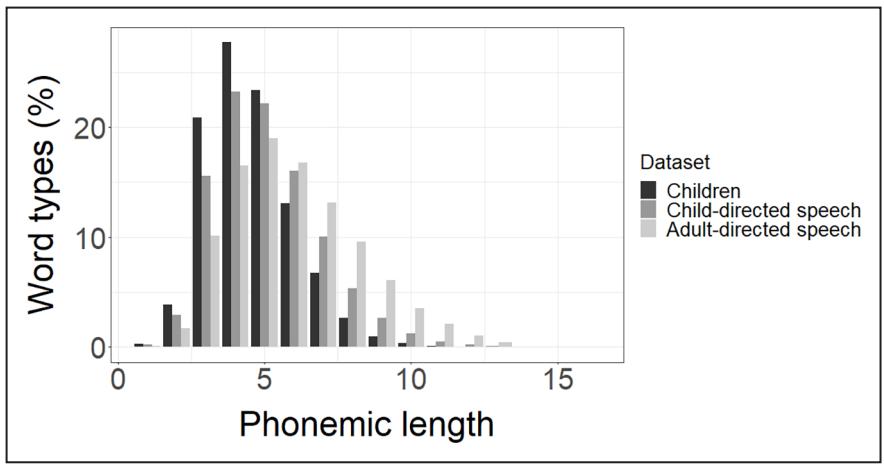

Figure 2. Percentage of Word Types Varying in Phonemic Length, for Each Sample Corpus (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech).

children's productive vocabularies (4051 phonological word types), we capture many more words than one would if using CDIs (~ 650 words); and (2) the number of word types increases with the age at which speech is directed (i.e. child-directed vs adult-directed).

Word length

As one may expect, child-directed speech is characterized by the use of shorter word types on average than adult-directed speech (see Figure 2). Table 2 also shows that child productive vocabularies contain proportionally more short word types than child-directed speech and in turn, child-directed speech contains proportionally more short word types than adult-directed speech.

Word frequency

Figure 3 shows how the proportion (plotted as density) of word types vary by word frequency across the three corpora when computing word frequency based on child-directed

Table 2. Kolmogorov–Smirnov comparisons for phonemic word length between different corpora, with 95% bootstrapped confidence intervals (2.5% and 97.5% percentiles, 1000 iterations).

Comparison	*D*	*p*	2.5%	97.5%
Children vs child-directed speech	.12	< .001	.10	.14
Child- vs adult-directed speech	.17	< .001	.15	.18
Children vs adult-directed speech	.29	< .001	.27	.31

The *p* values and confidence intervals were corrected using Holm's correction.

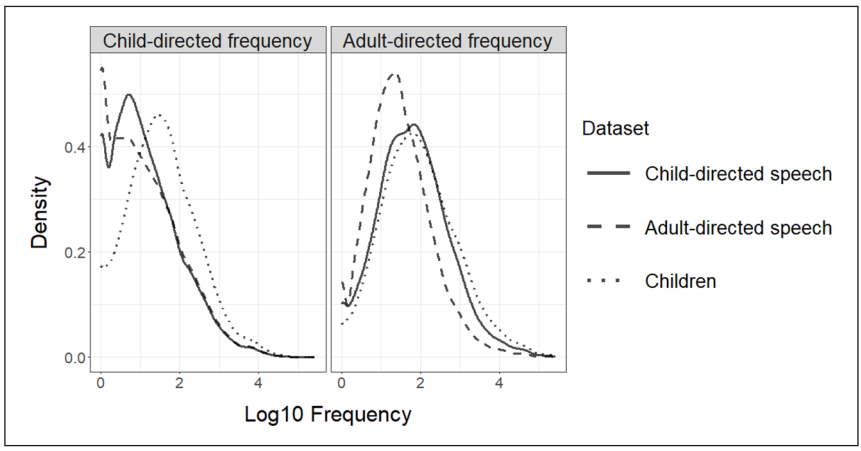

Figure 3. Density Plot Showing How Often Words Occur across the Frequency Range in Each Sample (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech), with Log10(Frequency) Computed Using the Child Corpus (Left Panel) and Adult Corpus (Right Panel). For Example, the Peak of the Dotted Line in the Right Panel Is at a Log10(Frequency) of Just Below 2, Indicating that the Words Children Produce Most Often Have a Log10(Frequency) of Just Below 2 When Using the Adult Corpus to Measure Frequency. Note that Because the Phonetic Dictionary Likely Omits Words That Rarely Occur, Frequencies Calculated from the Adult Corpora In Particular Do Not Contain Many Words Having Very Low Frequencies.

speech and adult-directed speech. Regardless of the corpora used to compute frequencies, children produce greater numbers of high frequency word types relative to those that appear in child-directed speech, and in turn, child-directed speech contains greater numbers of high frequency word types than adult-directed speech. These effects are confirmed by Kolmogorov–Smirnov comparisons (see Table 3). Of note is that the density plots in Figure 3 show that the frequencies of words used by children and in child-directed speech are much closer to one another than those involving adult-directed speech – regardless of which corpora are used to compute word frequency – suggesting that children's productive vocabularies are influenced by the frequency of the words they hear in child-directed speech.[3]

Table 3. Kolmogorov–Smirnov comparisons for Log10 word frequency between different corpora, with 95% bootstrapped confidence intervals (2.5% and 97.5% percentiles, 1000 iterations).

Comparison	Word frequency measure	*D*	*p*	2.5%	97.5%
Children vs child-directed speech	Child-directed	.23	< .001	.21	.26
Child- vs adult-directed speech	Child-directed	.05	< .001	.04	.07
Children vs adult-directed speech	Child-directed	.24	< .001	.22	.27
Children vs child-directed speech	Adult-directed	.06	< .001	.05	.09
Child- vs adult-directed speech	Adult-directed	.18	< .001	.17	.20
Children vs adult-directed speech	Adult-directed	.23	< .001	.22	.26

Comparisons are performed when computing Log10 word frequency using child-directed speech and adult-directed speech.

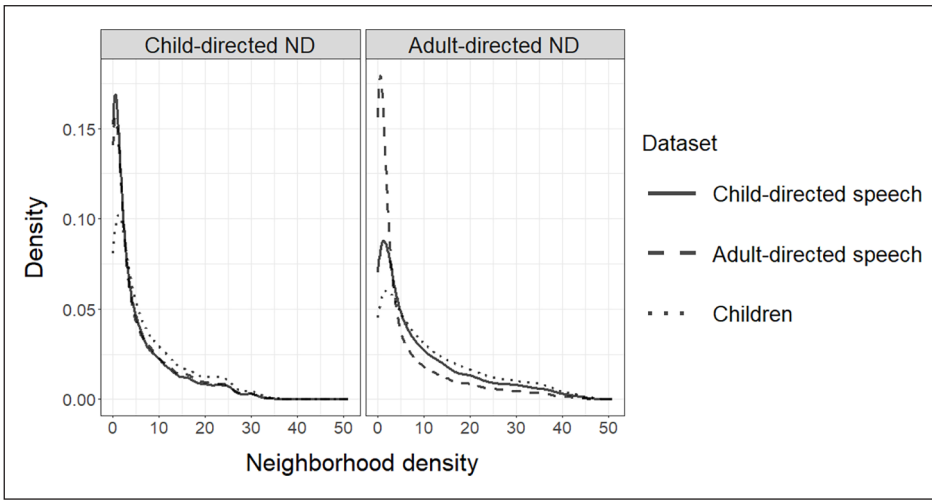

Figure 4. Density Plot Showing How Often Words Occur across the Neighborhood Density Range in Each Sample (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech), with Neighborhood Density Computed Using the Child Corpus (Left Panel) and Adult Corpus (Right Panel). Regardless of the Sample Used, the Peaks Are Always Just Above 0, Indicating that a Low Neighborhood Density Is Most Frequent across All Samples.

Neighborhood density

Figure 4 shows how the proportion (plotted as density) of words produced by children and used in child-directed and adult-directed speech vary in neighborhood density, when calculating neighborhoods using child-directed speech and adult-directed speech. In general, child productive vocabularies contain word types from more dense neighborhoods than child-directed speech (i.e. lower proportions of low neighborhood density words for child productive vocabularies relative to child-directed speech but higher proportions of higher density neighborhood words), and child-directed speech contains word types from more

Table 4. Kolmogorov–Smirnov comparisons for neighborhood density between different corpora, with 95% bootstrapped confidence intervals (2.5% and 97.5% percentiles, 1000 iterations).

Comparison	Neighborhood density measure	*D*	*p*	2.5%	97.5%
Children vs child-directed speech	Child-directed	.13	< .001	.11	.16
Child- vs adult-directed speech	Child-directed	.02	.083	.01	.04
Children vs adult-directed speech	Child-directed	.13	< .001	.10	.15
Children vs child-directed speech	Adult-directed	.11	< .001	.09	.14
Child- vs adult-directed speech	Adult-directed	.19	< .001	.17	.21
Children vs adult-directed speech	Adult-directed	.30	< .001	.27	.32

Comparisons are performed when computing neighborhood density using child-directed speech and adult-directed speech.

dense neighborhoods than adult-directed speech (see also, statistical comparisons, Table 4). However, Figure 4 clearly shows at least two things of note: (1) differences across children, child-directed speech, and adult-directed speech mainly stem from differences in the quantities of low neighborhood words in children's productive vocabularies (and slightly lower amounts of low neighborhood words in child-directed speech); and (2) differences between child-directed speech and adult-directed speech are more pronounced when neighborhood is calculated using the adult-directed corpora. At this stage it is not clear whether the neighborhood densities of words used in child-directed speech influence child vocabularies relative to adult-directed speech but we will return to this when examining predictors of children's productive vocabulary.

Phonotactic probability

Figure 5 shows how the proportion (plotted as density) of word types vary by phonotactic probability for children, child-directed speech, and adult-directed speech, for each corpus used to calculate phonotactic probability. The figure shows that children produce greater numbers of low phonotactic probability words relative to those that appear in child-directed speech, and in turn, child-directed speech contains greater numbers of low phonotactic probability word types than adult-directed speech. These conclusions are generally confirmed by Kolmogorov–Smirnov and confidence interval analyses (see Table 5). That said, one can see from Figure 5 that any differences are smaller when phonotactic probability is computed using child-directed speech than it is when computed using adult-directed speech. As with neighborhood density, we will examine this further when analyzing predictors of children's productive vocabularies.

Up to this point, we have examined our four measures independently, showing how each measure differs between child-directed speech and adult-directed speech. However, these measures interact with one another. For example, as words get longer, both their frequency of occurrence and neighborhood density decline (see Jones et al., 2021 for discussion). In addition, while we have examined whether there are differences in the measures across child-directed and adult-directed speech, we do not know if these differences affect children's vocabulary learning. In the next section, we therefore examine

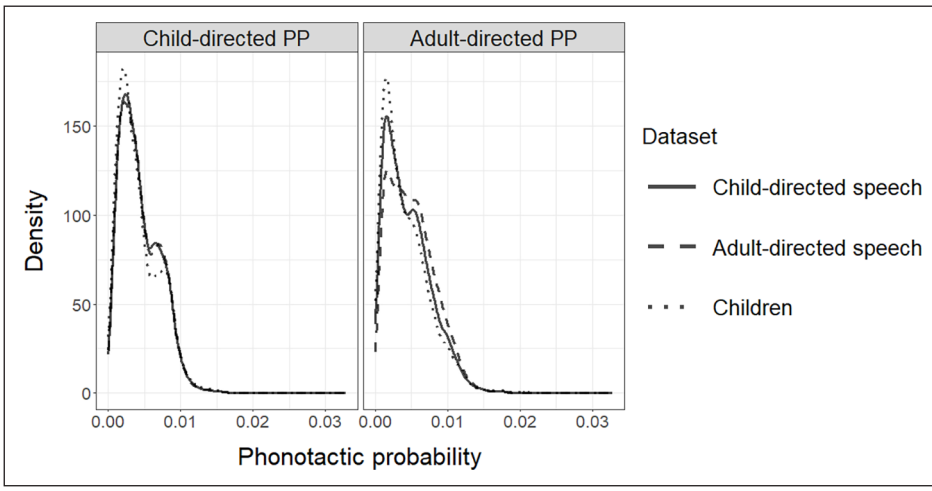

Figure 5. Density Plot Showing How Often Words Occur Across the Phonotactic Probability Range in Each Sample (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech), with Phonotactic Probability Computed Using the Child Corpus (Left Panel) and Adult Corpus (Right Panel). Regardless of the Sample Used, the Peaks Are Always Just Above 0, Indicating that a Low Phonotactic Probability Is Most Frequent across All Samples.

Table 5. Kolmogorov–Smirnov comparisons for phonotactic probability between different corpora, with 95% bootstrapped confidence intervals (2.5% and 97.5% percentiles, 1000 iterations).

Comparison	Phonotactic probability measure	*D*	*p*	2.5%	97.5%
Children vs child-directed speech	Child-directed	.05	<.001	.04	.08
Child- vs adult-directed speech	Child-directed	.01	.977	.01	.03
Children vs adult-directed speech	Child-directed	.05	<.001	.03	.08
Children vs child-directed speech	Adult-directed	.06	<.001	.04	.08
Child- vs adult-directed speech	Adult-directed	.09	<.001	.07	.11
Children vs adult-directed speech	Adult-directed	.14	<.001	.12	.16

Comparisons are performed when computing phonotactic probability using child-directed speech and adult-directed speech.

whether the measures – together with pairwise interactions – are predictive of children's productive vocabularies.

Word length, word frequency, neighborhood density, and phonotactic probability as predictors of child productive vocabularies

We now examine whether our four measures and their pairwise interactions are predictive of child productive vocabularies on the basis of which sample the measures are computed from, that is, whether child-directed speech is more predictive of child productive

Table 6. Number of nouns that are produced by children together with those that appear in maternal utterances but not in the child utterances (i.e. not produced by children), by phonemic length.

Phonemic length	Produced	Not produced	Ratio
2	335	95	3.53
3	2423	1092	2.22
4	2261	1426	1.59
5	1666	1287	1.29
6	725	808	0.90
7	449	588	0.76
8	195	287	0.68
9	52	128	0.41
10	13	79	0.16
11	5	29	0.17

vocabularies than adult-directed speech. In order to do this, we need to see whether there is anything particular in relation to our four measures regarding words that the children produce versus words that they fail to produce. We therefore capitalized on the Manchester corpus containing both maternal and child utterances that are part of ongoing caregiver–child conversations. To make our analyses tractable, we only considered nouns. All nouns appearing in child utterances were classified as *produced* nouns. All nouns appearing in maternal utterances but not in the child utterances were classified as *not produced*, since we have no evidence of any of the children speaking those nouns. A major advantage of using the maternal utterances that are part of the ongoing mother–child conversations is that all nouns have been heard by the children, regardless of whether they subsequently produced the nouns. This method also allows us to examine whether there are any particular features (across our four measures and their pairwise interactions) of produced nouns that distinguish them from those that are not produced. Table 6 shows the sample size of produced and not produced nouns, by phonemic length.[4] Produced nouns outweigh not produced nouns for words up to five phonemes in length, but beyond that there are more nouns that children fail to produce than those that they produce. In subsequent analyses, nouns of length greater than eight phonemes were discarded due to low sample sizes.

Figure 6 shows the distribution of each of the four measures used in the current study, for produced and not produced nouns when the measures are computed using child-directed speech and adult-directed speech. The figure largely confirms our previous findings whereby children's productive vocabularies are dominated by short words, high frequency words, low phonotactic probability words and proportionally greater numbers of words from dense neighborhoods – regardless of the source used to compute the measures. That said, the evidence for phonotactic probability is marginal. The data also show that all of the measures are skewed. For analysis, we transformed (standardized) the data so that odds ratios can be directly compared. Inspection of the model's residuals using the DHARMa package (Hartig, 2022) indicated no problems relating to deviation from the expected distribution of the data.

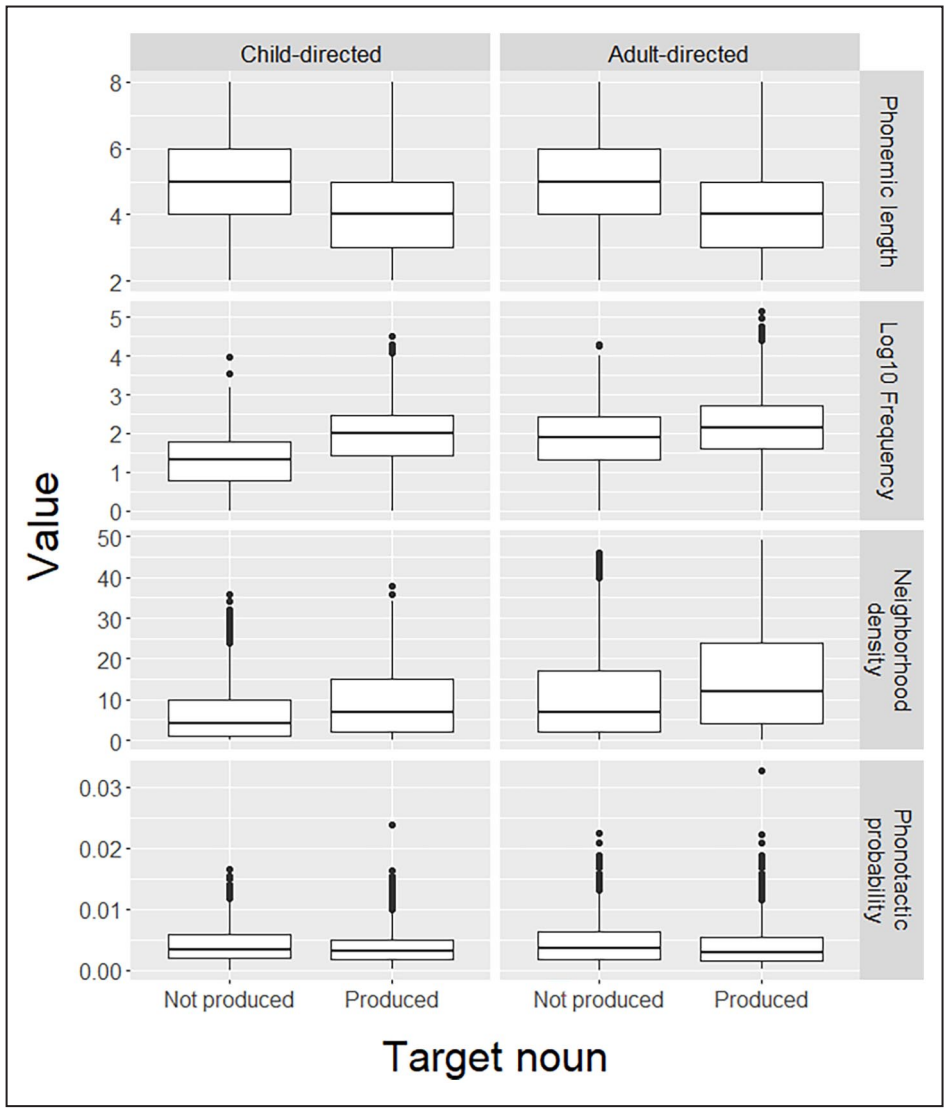

Figure 6. Box Plot Comparisons of Nouns Produced by Children in the Manchester Corpus (Produced) and Nouns Appearing in the Maternal Utterances of the Corpus That Were Not Produced by Children (not Produced) for Each of the Four Measures That Are the Focus of the Current Study, with Each Measure Computed Using the Child-Directed Speech and Adult-Directed Speech Corpora.

Table 7 provides mixed effects logistic regression analyses using our four measures as predictors of whether a noun was produced or not, when computing the measures using the child-directed speech and adult-directed speech corpora. There are at least three notable aspects of this analysis that favor child-directed speech over adult-directed speech.

First, child-directed speech explains over twice the variability in child productive vocabulary than adult-directed speech (conditional R^2 of .263 vs .124). Second, although both child-directed and adult-directed speech show word length and word frequency as significant predictors (with neighborhood density and phonotactic probability notably absent), the odds ratios[5] indicate the effect of word length is similar across the two different corpora, whereas word frequency has a much greater influence when measured by child-directed speech than adult-directed speech. Third, there is an additional interaction

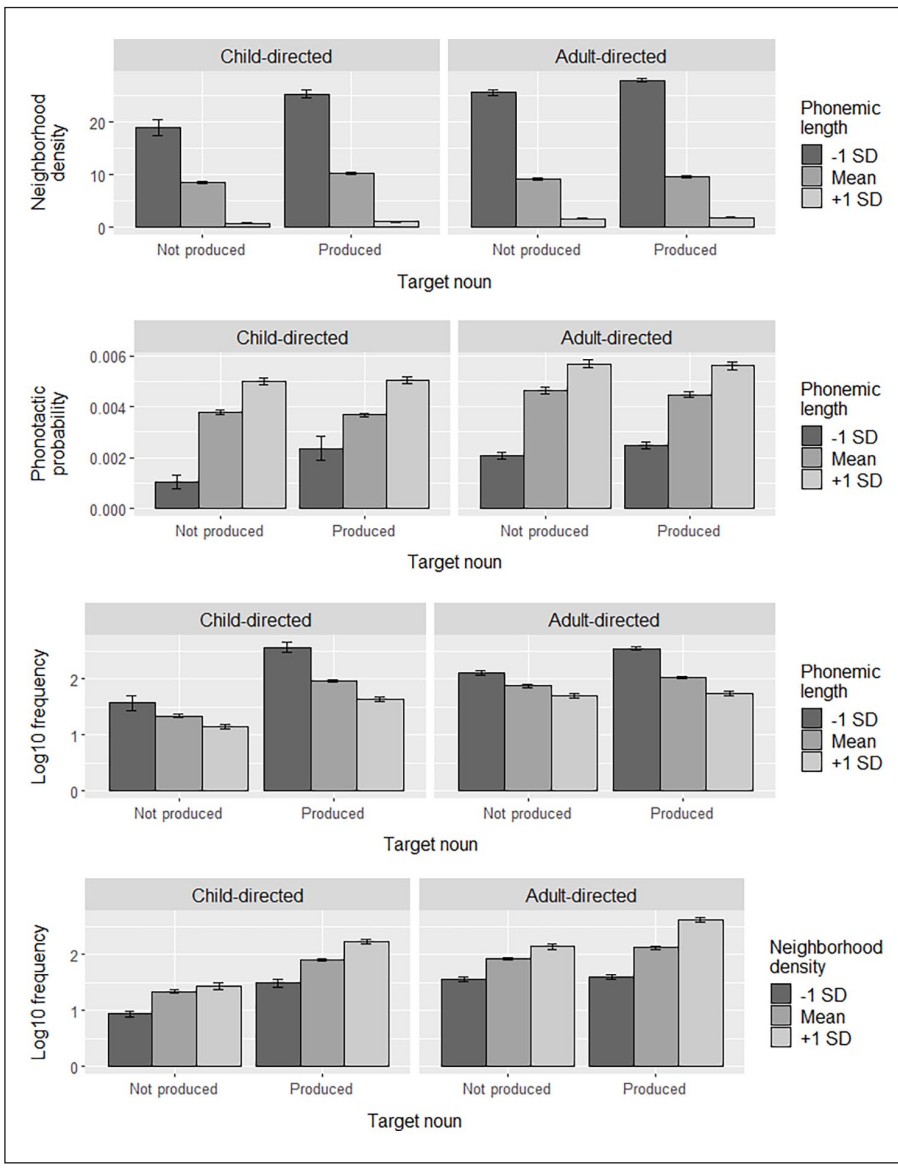

Figure 7. *(Continued)*

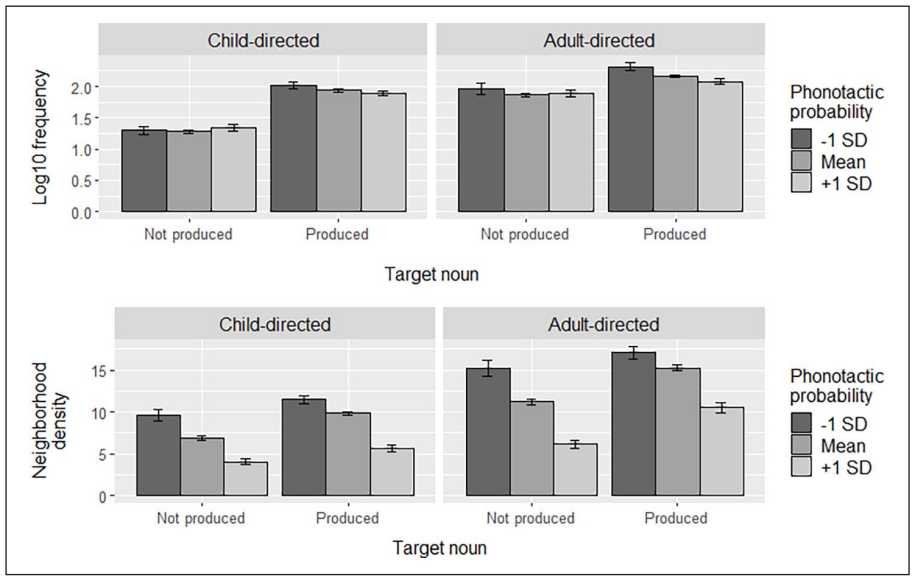

Figure 7. Box Plot Comparisons of Pairwise Interactions across Word Length, Log10 Frequency, Neighborhood Density, and Phonotactic Probability, When These Measures Are Computed Using Either Child-Directed Speech or Adult Directed Speech. Error Bars Indicate Confidence Intervals.

effect for child-directed speech compared with adult-directed speech. For completeness, Figure 7 investigates all pairwise interaction effects, though the key interactions from Table 7 are those between neighborhood density and frequency and neighborhood density and phonotactic probability. The most influential based on odds ratios is between frequency and neighborhood density, showing that while frequency influences child noun production, it is particularly so when the nouns in question are also high in neighborhood density. The neighborhood density and phonotactic probability interaction shows that as phonotactic probability increases, neighborhood density has an increasingly influential effect: children are much more likely to produce nouns having a high phonotactic probability when those nouns are high in neighborhood density. In short, the interactions show that neighborhood density is an important factor, particularly when measured using child-directed speech.

Discussion

This article set out to examine whether the lexical characteristics of children's early vocabularies might be driven by the lexical characteristics of the input they receive. We compared large-scale samples of child-directed speech to the same quantity of adult-directed speech, finding that a plausible reason that children's vocabularies contain short words, high frequency words, words that are low in phonotactic probability, and words that are proportionally high in neighborhood density, is that child-directed speech

Table 7. Mixed-effects logistic regression analyses predicting word production of nouns (levels: produced/not produced) as a function of the four continuous predictors (phonemic length, Log10 word frequency, neighborhood density, and phonotactic probability), and the pairwise interactions of those predictors.

Predictors	Child-directed model			Adult-directed model		
	Odds ratios	*CI*	*p*	*Odds ratios*	*CI*	*p*
(Intercept)	1.78	1.43–2.22	**<.001**	1.52	1.22–1.89	**<.001**
Length	0.83	0.75–0.93	**.001**	0.78	0.71–0.86	**<.001**
Frequency	2.47	2.35–2.59	**<.001**	1.34	1.28–1.39	**<.001**
Neighborhood	1.03	0.91–1.16	.654	1.08	0.98–1.19	.134
Phonotactic	0.99	0.94–1.03	.556	1.01	0.97–1.05	.683
Length × neighborhood	0.95	0.86–1.04	.242	1.00	0.92–1.07	.904
Length × phonotactic	0.98	0.92–1.05	.586	0.97	0.91–1.03	.290
Length × frequency	1.01	0.94–1.09	.727	1.00	0.94–1.07	.988
Frequency × neighborhood	1.16	1.08–1.25	**<.001**	1.23	1.15–1.31	**<.001**
Frequency × phonotactic	0.97	0.92–1.02	.237	0.96	0.92–1.00	.070
Neighborhood × phonotactic	1.09	1.02–1.17	**.015**	1.06	0.99–1.13	.076
Random effects						
σ^2	3.29			3.29		
τ_{00}	$.13_{id}$			$.14_{id}$		
ICC	.04			.04		
N	12_{id}			12_{id}		
Observations	12,299			13,328		
Marginal R^2/conditional R^2	.234/.263			.087/.124		

Neighborhood = neighborhood density; phonotactic = phonotactic probability. CI: confidence interval; ICC: intraclass correlation coefficient.
Random intercept of child is included. Child predictors are computed using standardized child-directed speech and adult-directed speech measures, so that odds ratios can be compared. Significant effects are highlighted in bold.

contains greater proportions of words that fit these categories than adult-directed speech. Regression analyses then examined the extent to which these factors – measured using either child-directed speech or adult-directed speech – were predictive of children's productive vocabularies. Child-directed speech explained over twice the variability in children's noun productions than did adult-directed speech, with the frequency of encounter of a noun in child-directed speech being a key factor and word length a secondary factor. However, there were also two interaction effects relating to how child-directed speech influenced children's productive vocabularies, both involving neighborhood density. The lexical characteristics of the words used in child-directed speech clearly play an important role in children's vocabulary acquisition. We will first consider the significant predictors relating to child-directed speech and their implications for children's vocabulary acquisition and then consider what the differences in lexical characteristics between child-directed speech and adult-directed speech may mean.

It is by no means surprising that frequency plays a vital role in what words children learn from child-directed speech. Swingley and Humphrey (2018) found frequency of occurrence in child-directed speech to be far and away the most consistent predictor of receptive and productive vocabularies in 15-month-old infants; child-directed speech is littered with repetitions that push frequency of occurrence upward (Hills, 2013); and quantity of speech in early childhood is a better predictor of vocabulary than the diversity of vocabulary (Jones & Rowland, 2017; Rowe, 2012). Our child-directed and adult-directed speech samples vastly differed in the number of phonological word types (8573 vs 29,539). Given the samples were matched for word tokens, this means greater frequency, and hence greater repetition, of many word types in child-directed speech relative to adult-directed speech, further supporting the aforementioned literature. Children also seem to find it easier to learn short words over long words, which is also not altogether surprising since almost every learning mechanism follows this pattern (e.g. PARSER: Perruchet & Vinter, 1998; TRACX: French et al., 2011). Yet frequency and word length cannot be the only key factors in early word learning. They have to give way to other developmental factors, not least because at the tail end of the developmental continuum – adult-directed speech – there is a large amount of infrequent words and long words (see Figures 2 and 3). This means that other lexical characteristics have to come into play at the expense of frequency and length as development progresses. The beginnings of this transition are arguably present in the current data.

There are two effects, both involving neighborhood density, that warrant some scrutiny. Neighborhood density influences the effects of both frequency and phonotactic probability, with frequent words and words containing biphone pairs that occur frequently being more likely to be produced by children if the words also have many similar sounding other words in child-directed speech. While there are numerous studies that have shown neighborhood density to be linked to vocabulary acquisition (e.g. Stokes, 2010; Storkel, 2009), there are others that show little or no effect (e.g. Swingley & Humphrey, 2018). One plausible explanation for this is that neighborhood effects may emerge later in development, with those studies showing an effect of neighborhood examining child ages of around 24 months and those showing no effect examining child ages of around 15 months. In addition, many studies examine receptive vocabulary; given that production lags behind comprehension (and production may also be more influenced by neighborhood than comprehension given there is an articulatory factor), our production data (involving an average child age of 28 months) are consistent with the hypothesis that neighborhood is an effect that appears at around 2 years of age.

However, one has to explain why neighborhood may influence vocabulary acquisition. In and of itself, the fact that child-directed speech contains proportionally more words with many neighbors than adult-directed speech does not tell us why children benefit. In order for children to realize such benefits, they have to know something about the neighboring words – so children must learn something about the neighbors themselves, heard in the child-directed input, and use this information to their advantage. One possibility (Walley, 1993; Walley et al., 2003) is that children initially represent words as unanalyzed wholes, but by hearing many similar sounding words, come to realize that they need to learn segmental detail for such words. Such segmental detail presumably facilitates the learning of other similar sounding, neighboring words. An alternative but not inconsistent view (Jones et al., 2021) is that children learn bottom-up from exposure

to linguistic input, with greater information being learned about the constituent sequences of sounds within a word as exposure to that word increases. On this view, children learn sound sequences within the words they hear. Should some of those words be neighbors of a to-be-learned word, then the overlap in learned phonological knowledge from neighbors of the new word will scaffold learning of the new word. The traditional definition of neighbor is therefore dispensed with since neighboring words are those that have some overlap in the sequence of sounds that comprise the words. The added advantage of Jones et al.'s (2021) view is that it also explains why both word length and word frequency are influential in children's word learning – since gradual learning is based on exposure to words, short words and frequent words are learned more quickly. In the former case, this is because short words require fewer exposures to be learned than do long words. In the latter case it is because frequency is directly linked to how quickly a word is learned, with greater exposure to the word leading to faster learning of it.

The hypothesis put forward by Jones et al. (2021) also makes sense of the interaction effects seen in the current data. Frequency effects are stronger when there are many words that share sequences of sounds compared with only a few; a given word is learned more quickly because there is both breadth of knowledge of the constituent phoneme sequences (from the many neighboring words) and the 'neighboring' knowledge is learned quickly (due to their high frequency). The interaction between phonotactic probability and neighborhood, whereby learning high phonotactic probability words is helped by the words having high neighborhoods, is also readily explained. Having frequent sound sequences that also appear across many words increases the chances of phonological knowledge being learned for those sound sequences, and in turn, this increases the chances of learning the word itself.

We now turn to the contributions of child-directed speech and adult-directed speech in children's vocabulary acquisition. One of the key findings of the current study is that there are differences between child-directed speech and adult-directed speech on every one of the four key measures used. This has implications for those studies that have examined children's word learning using measures that are derived from adult-directed speech or adult-directed written literature. If child-directed speech differs from adult-directed speech on all measures, one can expect similar differences for adult-directed written literature. Although previous studies using adult-directed written literature to estimate word frequency, phonotactic probability, and neighborhood density (e.g. Stokes, 2010; Storkel, 2009) are broadly consistent with the results presented here when using child-directed speech corpora, they are likely under-estimating the influence of these measures.

Our data show that there are key aspects of child-directed speech – word frequency, word length, and neighborhood density – that influence child productive vocabularies. In addition, the influence of these variables is much stronger than for adult-directed speech. This gives some explanation as to why overheard speech from adults does not significantly influence child vocabularies (Weisleder & Fernald, 2013).[6] A natural assumption from this data is that greater repetitions of words and an increasing number of short words will boost children's vocabulary growth. However, a clear caveat is the emerging effect of neighborhood density discussed earlier. Frequency and length have to give way to other variables given that with age, children's vocabularies involve increasing amounts of infrequent and long words. A plausible factor, therefore, is neighborhood – and by providing children with a linguistic input

that is littered with neighboring words (or overlapping sound sequences), their vocabulary acquisition is likely to be scaffolded based on their knowledge of neighboring words.

A final note concerns the fundamental differences between child-directed speech and adult-directed speech that alter the lexical measures used in this article. Adult-directed speech uses longer words than child-directed speech (Figure 2), and also a far greater number of word types (Table 1) that reduce word frequency (Figure 3, right panel). In all likelihood, these will change neighborhood density, due to the presence of more word types, and phonotactic probability, due to changes in word frequencies. Would our lexical measures be similar across child-directed speech and adult-directed speech if the adult-directed speech matched the child-directed speech on word length and number of word types? Our data suggest that this may well be the case. When our lexical measures were computed from child corpora (e.g. neighborhood density and phonotactic probability being computed from words directed to children) there were no differences between child-directed speech and adult-directed speech for neighborhood density and phonotactic probability (Table 4, $p = .083$ and Table 5, $p = .977$, respectively).

In summary, this article has compared large-scale child-directed and adult-directed speech to investigate whether the lexical characteristics of words in child-directed speech differ from those of adult-directed speech in ways that influence the lexical properties of words that appear in child productive vocabularies. We found the lexical properties of child vocabularies are predicted by three of the four lexical characteristics under examination, either directly or indirectly. Children learn more short than long words, more frequent than infrequent words, and more neighboring words than words with few neighbors. These characteristics are significantly exaggerated in child-directed speech relative to adult-directed speech. Furthermore, the predictive potency of the measures change as the corpora used to measure them change, with child-directed speech explaining over twice the variability in child productive vocabularies than adult-directed speech. Children's word learning is clearly influenced by the lexical characteristics of the words they hear, but researchers need to be aware that using different corpora to calculate lexical measures may influence their perceived effect on children's word learning.

Author contributions

Gary Jones: Conceptualization; Formal analysis; Investigation; Methodology; Project administration; Writing – original draft; Writing – review & editing.
Francesco Cabiddu: Data curation; Formal analysis; Methodology; Resources; Software; Writing – review & editing.
Doug J. K. Barrett: Formal analysis; Writing – review & editing.
Antonio Castro: Formal analysis; Writing – review & editing.
Bethany Lee: Data curation; Investigation; Writing – review & editing.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Gary Jones https://orcid.org/0000-0003-3867-9947
Francesco Cabiddu https://orcid.org/0000-0001-9692-4897

Notes

1. This apparent contradiction is likely related to changes in the interaction taking place (Dong et al., 2021; see also Golinkoff et al., 2015).
2. Phonemic or syllabic word length (i.e. we do not consider spoken duration).
3. Note that the number of word types within each corpus directly influences word frequency: adult-directed speech has many more word types than the children and child-directed speech (see Table 1), which inevitably means that adult-directed speech will have greater numbers of low frequency words.
4. Note that our analyses examine lexical characteristics but also the variance explained by child, and so the produced nouns in Table 6 are totals across the 12 children in the corpus (e.g. for 'truck', the produced nouns column would include a score for every child who produced 'truck'). This explains why the produced noun totals are greater than the 4051 words in Table 1, which looks at the corpora as a whole and are not restricted to nouns.
5. For odds ratios, the greater the deviation from 1.00, the more the predictor wields an influence on the outcome measure.
6. Note there are at least three possible alternative explanations for why adult-directed speech may be less predictive of children's vocabulary learning than child-directed speech. First, child-directed speech is more geared toward gaining the child's attention (e.g. differences in intonation). Second, it may be the case that the language used in child-directed speech fits more with the child's own semantic networks (e.g. see Hills, 2013). Third, child-directed speech is produced when children are engaged in interaction, so there may be additional non-verbal cues that benefit word learning.

References

Baddeley, A. D., Thomson, N., & Buchanan, M. (1975). Word length and the structure of short-term memory. *Journal of Verbal Learning and Verbal Behavior*, *14*, 575–589.

Bernstein Ratner, N., & Rooney, B. (2001). How accessible is the lexicon in motherese? In J. Weissenborn, & B. Hohle (Eds.), *Approaches to bootstrapping: Phonological, lexical, syntactic and neurophysiological aspects of early language acquisition* (Vol.1, pp. 71–78.) Benjamins Publishing Company.

British National Corpus. (2007). *The British National Corpus* (Version 3, BNC XML Edition) [On-line Database]. http://www.natcorp.ox.ac.uk/

Brent, M. R., & Siskind, J. M. (2001). The role of exposure to isolated words in early vocabulary development. *Cognition*, *81*, B33–B44.

Carlson, M. T., Sondregger, M., & Bane, M. (2014). How children explore the phonological network in child-directed speech: A survival analysis of children's first word productions. *Journal of Memory and Language*, *75*, 159–180.

Clark, E. V., Gelman, S. A., & Lane, N. M. (1985). Compound nouns and category structure in young children. *Child Development*, *56*, 84–94.

Coady, J. A., & Aslin, R. N. (2003). Phonological neighbourhoods in the developing lexicon. *Journal of Child Language*, *30*, 441–469.

Dong, S., Gu, Y., & Vigliocco, G. (2021). The impact of child-directed language on children's lexical development. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *43*, 1444–1450.

Fernald, A. (1989). Intonation and communicative intent in mothers' speech to infants: Is the melody the message? *Child Development*, *60*, 1497–1510.

Forrester, M. (2002). Appropriating cultural conceptions of childhood: Participation in conversation. *Childhood*, *9*, 255–278.

French, R. M., Addyman, C., & Mareschal, D. (2011). TRACX: A recognition-based connectionist framework for sequence segmentation and chunk extraction. *Psychological Review*, *118*, 614–636.

Furrow, D., Nelson, K., & Benedict, H. (1979). Mothers' speech to children and syntactic development: Some simple relationships. *Journal of Child Language*, *6*, 423–442.

Gathercole, S. E., & Baddeley, A. D. (1989). Evaluation of the role of phonological STM in the development of vocabulary in children: A longitudinal study. *Journal of Memory and Language*, *28*, 200–213.

Golinkoff, R. M., Can, D. D., Soderstrom, M., & Hirsh-Pasek, K. (2015). (Baby)Talk to me: The social context of infant-directed speech and its effects on early language acquisition. *Current Directions in Psychological Science*, *24*, 339–344.

Goodman, J. C., Dale, P. S., & Li, P. (2008). Does frequency count? Parental input and the acquisition of vocabulary. *Journal of Child Language*, *35*, 515–531.

Hartig, F. (2022). *DHARMa: Residual diagnostics for hierarchical (multi-level / mixed) regression models* (R package version 0.4.5). https://CRAN.R-project.org/package=DHARMa

Hayes, D. P., & Ahrens, M. G. (1988). Vocabulary simplification for children: A special case of motherese? *Journal of Child Language*, *15*, 395–410.

Henry, A. (1995). *Belfast English and standard English: Dialect variation and parameter setting*. Oxford University Press.

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

Hoff, E., & Naigles, L. R. (2002). How children use input to acquire a lexicon. *Child Development*, *73*, 418–433.

Huttenlocher, J., Haight, W., Bryk, A., Seltzer, M., & Lyons, T. (1991). Early vocabulary growth: Relation to language input and gender. *Developmental Psychology*, *27*, 236–248.

Jones, G., Cabiddu, F., Andrews, M., & Rowland, C. (2021). Chunks of phonological knowledge play a significant role in children's word learning and explain effects of neighborhood size, phonotactic probability, word frequency and word length. *Journal of Memory and Language*, *119*, 104232.

Jones, G., & Rowland, C. F. (2017). Diversity not quantity in caregiver speech: Using computational modeling to isolate the effects of the quantity and the diversity of the input on vocabulary growth. *Cognitive Psychology*, *98*, 1–21.

Jones, G., Tamburelli, M., Watson, S. E., Gobet, F., & Pine, J. M. (2010). Lexicality and frequency in specific language impairment: Accuracy and error data from two nonword repetition tests. *Journal of Speech Language and Hearing Research*, *53*, 1642–1655.

Jones, G., & Witherstone, H. L. (2011). Lexical and sublexical knowledge influences the encoding, storage, and articulation of nonwords. *Memory & Cognition*, *39*, 588–599.

Lieven, E., Salomo, D., & Tomasello, M. (2009). Two-year-old children's production of multiword utterances: A usage-based analysis. *Cognitive Linguistics*, *20*, 481–507.

Liu, H. M., Kuhl, P. K., & Tsao, F. M. (2003). An association between mothers' speech clarity and infants' speech discrimination skills. *Developmental Science*, *6*, F1–F10.

MacWhinney, B. (2000). *The CHILDES project: The database* (Vol.2). Psychology Press.

Maekawa, J., & Storkel, H. L. (2006). Individual differences in the influence of phonological characteristics on expressive vocabulary development by young children. *Journal of Child Language*, *33*, 439–459.

McMurray, B., Kovack-Lesh, K. A., Goodwin, D., & McEchron, W. (2013). Infant directed speech and the development of speech perception: Enhancing development or an unintended consequence? *Cognition*, *129*, 362–378.

Montag, J. L., Jones, M. N., & Smith, L. B. (2018). Quantity and diversity: Simulating early word learning environments. *Cognitive Science*, *42*, 375–412.

Naigles, L. R., & Hoff-Ginsberg, E. (1998). Why are some verbs learned before other verbs? Effects of input frequency and structure on children's early verb use. *Journal of Child Language*, *25*, 95–120.

Perruchet, P., & Vinter, A. (1998). PARSER: A model for word segmentation. *Journal of Memory and Language*, *39*, 246–263.

R Core Team. (2018). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. https://www.R-project.org/

Rodriguez-Cuadrado, S., Baus, C., & Costa, A. (2018). Foreigner talk through word reduction in native/non-native spoken interactions. *Bilingualism*, *21*, 419–426.

Rowe, M. L. (2012). A longitudinal investigation of the role of quantity and quality of child-directed speech vocabulary development. *Child Development*, *83*, 1762–1774.

Rowland, C. F., & Fletcher, S. L. (2006). The effect of sampling on estimates of lexical specificity and error rates. *Journal of Child Language*, *33*, 859–877.

Roy, B. C., Frank, M. C., & Roy, D. (2009). Exploring word learning in a high-density longitudinal corpus. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*. Amsterdam, The Netherlands. https://www.media.mit.edu/cogmac/publications/roy-frank-roy-cogsci2009.pdf

Saxton, M. (2009). The inevitability of child directed speech. In S. Foster-Cohen (Ed.), *Language acquisition* (pp. 62–86). Springer.

Stokes, S. F. (2010). Neighborhood density and word frequency predict vocabulary size in toddlers. *Journal of Speech, Language, and Hearing Research*, *53*, 670–683.

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*, 291–321.

Storkel, H. L., & Lee, S. (2011). The independent effects of phonotactic probability and neighborhood density on lexical acquisition by preschool children. *Language and Cognitive Processes*, *26*, 191–211.

Suttora, C., Salerni, N., Zanchi, P., Zampini, L., Spinelli, M., & Fasolo, M. (2017). Relationships between structural and acoustic properties of maternal talk and children's early word recognition. *First Language*, *37*, 612–629.

Swingley, D., & Humphrey, C. (2018). Quantitative linguistic predictors of infants' learning of specific English words. *Child Development*, *89*, 1247–1267.

Theakston, A. L., Lieven, E. V., Pine, J. M., & Rowland, C. F. (2001). The role of performance limitations in the acquisition of verb-argument structure: An alternative account. *Journal of Child Language*, *28*, 127–152.

Thiessen, E. D., Hill, E. A., & Saffran, J. R. (2005). Infant-directed speech facilitates word segmentation. *Infancy*, *7*, 53–71.

Tommerdahl, J., & Kilpatrick, C. (2013). Analyzing reliability of grammatical production in spontaneous samples of varying length. *Journal of Child Language Teaching and Therapy*, *29*, 171–183.

Vaden, K. I., Halpin, H. R., & Hickok, G. S. (2009). Irvine phonotactic online dictionary, version 2. http://www.iphod.com

Walley, A. C. (1993). The role of vocabulary development in children's spoken word recognition and segmentation ability. *Developmental Review*, *13*, 286–350.

Walley, A. C., Metsala, J. L., & Victoria, M. (2003). Spoken vocabulary growth: Its role in the development of phoneme awareness and early reading ability. *Reading and Writing: An Interdisciplinary Journal*, *16*, 5–20.

Weisleder, A., & Fernald, A. (2013). Talking to children matters: Early language experience strengthens processing and builds vocabulary. *Psychological Science*, *24*, 2143–2152.

Wells, C. G. (1981). *Learning through interaction: The study of language development*. Cambridge University Press.

You, G., Bickel, B., Daum, M. M., & Stoll, S. (2021). Child-directed speech is optimized for syntax-free semantic inference. *Scientific Reports*, *11*, 1–11.

Appendix I

Plots and analyses when excluding all compound nouns

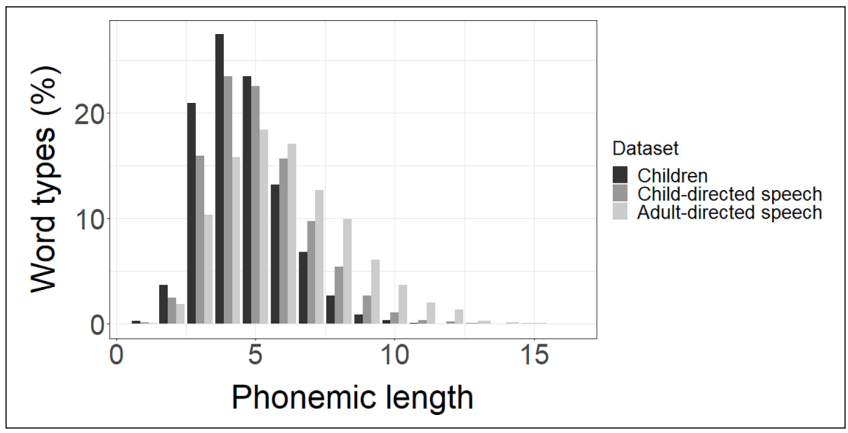

Figure 8. Percentage of Word Types Varying in Phonemic Length, for each Sample Corpus (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech) When Excluding All Compound Nouns.

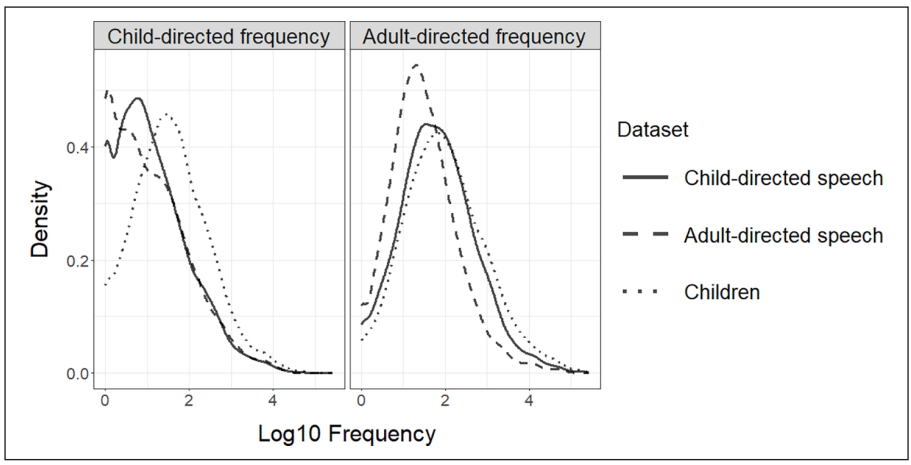

Figure 9. Density Plot of Word Types (Excluding Compound Nouns) Varying in Log10 Word Frequency, for Each Sample Corpus (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech) and with Word Frequency Computed Using Child-Directed Speech (Left Panel) and Adult-Directed Speech (Right Panel).

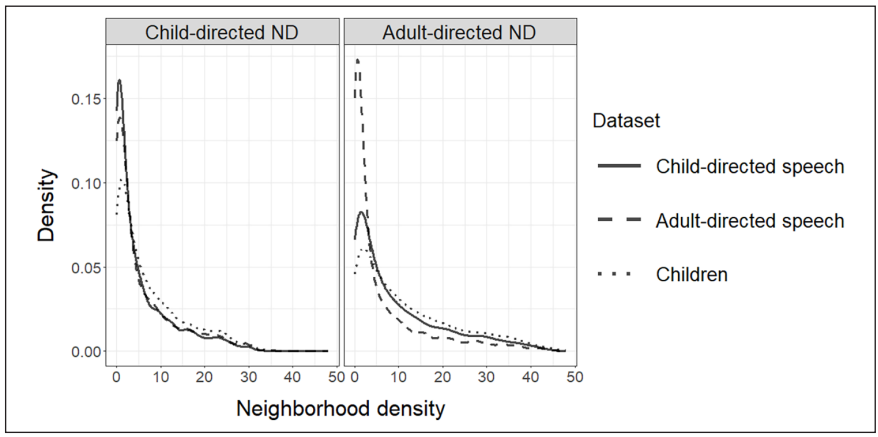

Figure 10. Density Plot of Word Types Varying in Neighborhood Density (Excluding Compound Nouns), for Each Sample Corpus (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech) and with Neighborhood Density Computed Using Child-Directed Speech (Left Panel) and Adult-Directed Speech (Right Panel).

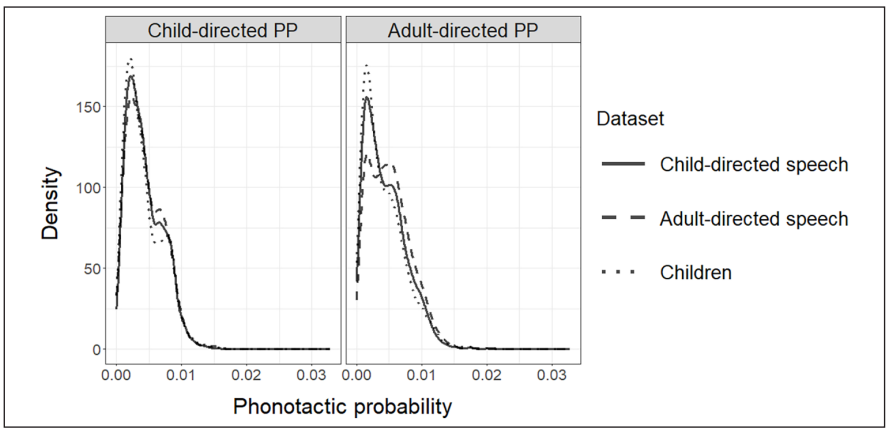

Figure 11. Density Plot of Word Types (Excluding Compound Nouns) Varying in Phonotactic Probability, for Each Sample Corpus (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech) and with Phonotactic Probability Computed Using Child-Directed Speech (Left Panel) and Adult-Directed Speech (Right Panel).

Table 8. Kolmogorov–Smirnov comparisons for phonemic word length between different corpora when excluding all compound nouns, with 95% bootstrapped confidence intervals (1000 iterations).

Comparison	*D*	*p*	Lower Bci	Upper Bci
Children vs child-directed speech	.11	<.001	.09	.13
Child- vs adult-directed speech	.18	<.001	.16	.2
Children vs adult-directed speech	.29	<.001	.27	.32

The *p* values and confidence intervals were corrected using Holm's correction.

Table 9. Kolmogorov–Smirnov comparisons for Log10 word frequency between different corpora (excluding compound nouns), with 95% bootstrapped confidence intervals (1000 iterations).

Comparison	Measure	*D*	*p*	Lower Bci	Upper Bci
Children vs child-directed speech	Child-directed	.25	<.001	.22	.27
Child- vs adult-directed speech	Child-directed	.06	<.001	.04	.09
Children vs adult-directed speech	Child-directed	.26	<.001	.23	.3
Children vs child-directed speech	Adult-directed	.07	<.001	.06	.1
Child- vs adult-directed speech	Adult-directed	.18	<.001	.16	.21
Children vs adult-directed speech	Adult-directed	.25	<.001	.22	.28

Comparisons are performed when computing Log10 word frequency using child-directed speech and adult-directed speech.

Table 10. Kolmogorov–Smirnov comparisons for neighborhood density between different corpora (excluding compound nouns), with 95% bootstrapped confidence intervals (1000 iterations).

Comparison	Measure	*D*	*p*	Lower Bci	Upper Bci
Children vs child-directed speech	Child-directed	.14	<.001	.11	.16
Child- vs adult-directed speech	Child-directed	.03	.455	.02	.05
Children vs adult-directed speech	Child-directed	.14	<.001	.1	.17
Children vs child-directed speech	Adult-directed	.1	<.001	.08	.13
Child- vs adult-directed speech	Adult-directed	.2	<.001	.18	.23
Children vs adult-directed speech	Adult-directed	.3	<.001	.27	.33

Comparisons are performed when computing neighborhood density using child-directed speech and adult-directed speech.

Table 11. Kolmogorov–Smirnov comparisons for phonotactic probability between different corpora (excluding compound nouns), with 95% bootstrapped confidence intervals (1000 iterations).

Comparison	Measure	*D*	*p*	Lower Bci	Upper Bci
Children vs child-directed speech	Child-directed	.05	.003	.03	.08
Child- vs adult-directed speech	Child-directed	.02	.614	.02	.06
Children vs adult-directed speech	Child-directed	.06	.001	.03	.1
Children vs child-directed speech	Adult-directed	.05	<.001	.03	.08
Child- vs adult-directed speech	Adult-directed	.1	<.001	.08	.13
Children vs adult-directed speech	Adult-directed	.15	<.001	.12	.18

Comparisons are performed when computing phonotactic probability using child-directed speech and adult-directed speech.

Appendix 2

Plots for individual word categories (noun, verb, adjective, and 'other')

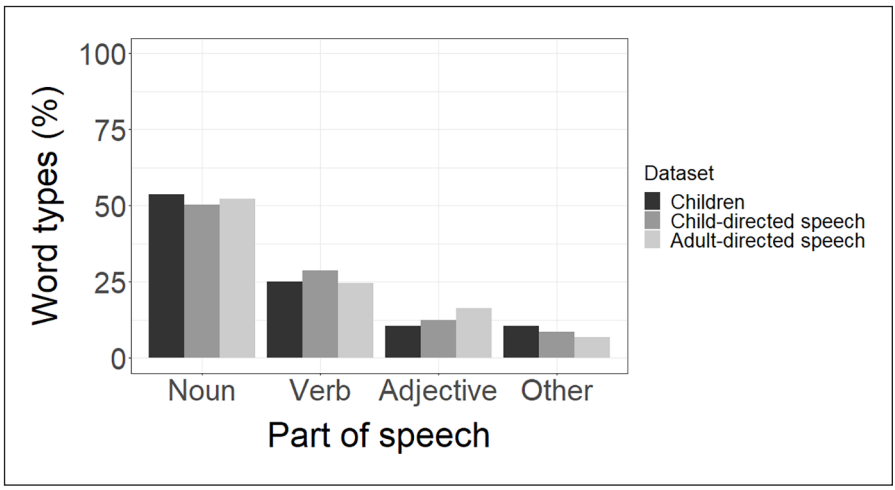

Figure 12. Distributions of Each Word Category within the Different Corpora (Children, Child-Directed Speech, Adult-Directed Speech).

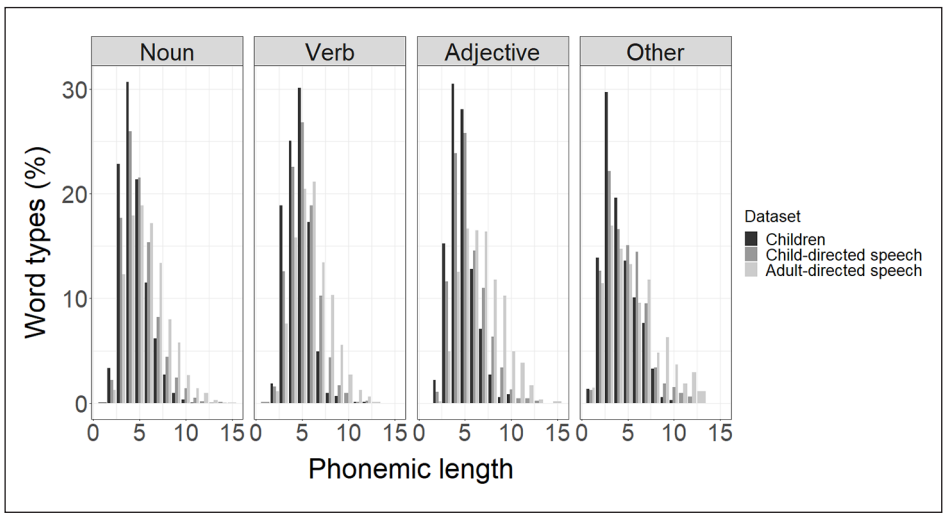

Figure 13. Distributions of Phonemic Length across Each Word Category within the Different Corpora (Children, Child-Directed Speech, Adult-Directed Speech).

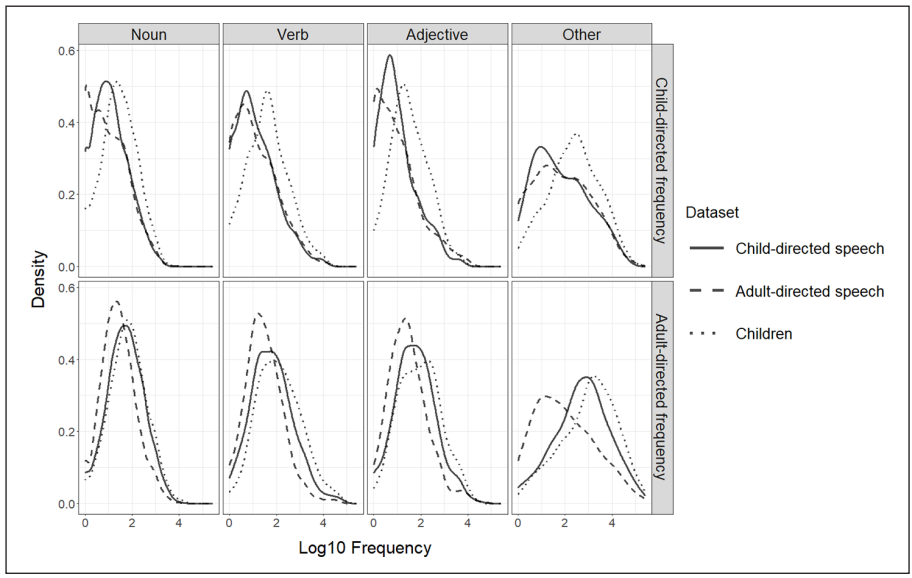

Figure 14. Density Plot of Word Types Varying in Log10 Word Frequency, for Each Word Category within Each Sample Corpus (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech).

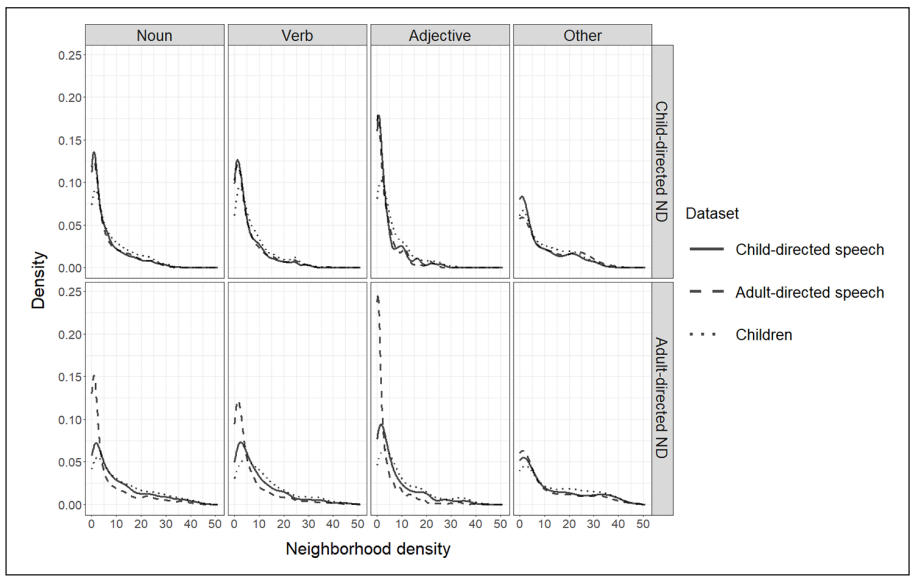

Figure 15. Density Plot of Word Types Varying in Neighborhood Density, for Each Word Category and Separated by Sample Corpus (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech).

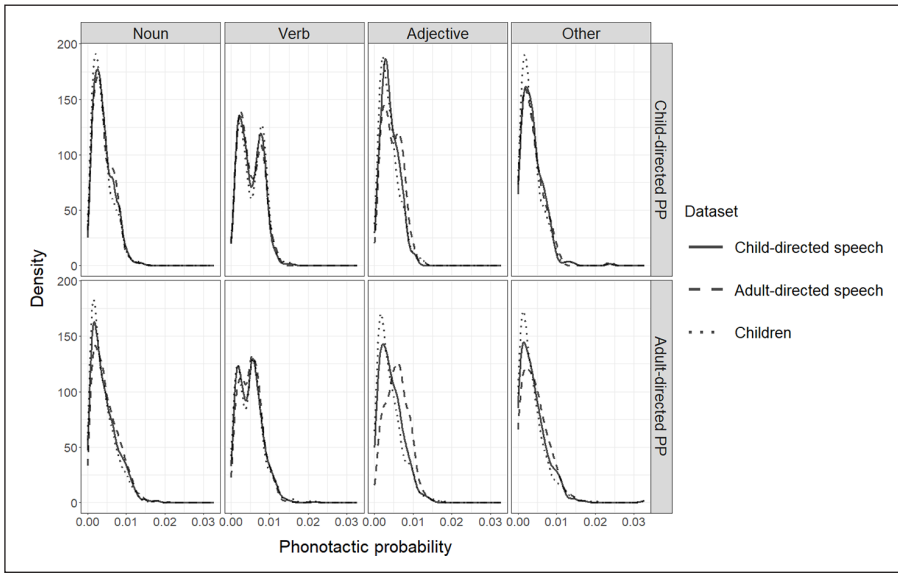

Figure 16. Density Plot of Word Types Varying in Phonotactic Probability, for Each Word Category and for Each Sample Corpus (Child Productive Vocabularies, Child-Directed Speech, Adult-Directed Speech).