An online calculator to compute phonotactic probability and neighborhood density on the basis of child corpora of spoken American English

HOLLY L. STORKEL AND JILL R. HOOVER

University of Kansas, Lawrence, Kansas

An online calculator was developed (www.bncdnet.ku.edu/cml/info_ccc.vi) to compute *phonotactic probability*—the likelihood of occurrence of a sound sequence—and *neighborhood density*—the number of phonologically similar words—on the basis of child corpora of American English (Kolson, 1960; Moe, Hopkins, & Rush, 1982) and to compare its results to those of an adult calculator. Phonotactic probability and neighborhood density were computed for a set of 380 nouns (Fenson et al., 1993) using both the child and adult corpora. The child and adult raw values were significantly correlated. However, significant differences were detected. Specifically, child phonotactic probability was higher than adult phonotactic probability, especially for high-probability words, and child neighborhood density was lower than adult neighborhood density, especially for words with high-density neighborhoods. These differences were reduced or eliminated when relative measures (i.e., *z* scores) were used. Suggestions are offered regarding which values to use in future research.

Many recent studies of spoken language processing by children have considered the role of *phonotactic probability*—the likelihood of occurrence of a sound (phoneme) sequence—and *neighborhood density*—the number of phonologically similar words—in word recognition (Garlock, Walley, & Metsala, 2001; Mainela-Arnold, Evans, & Coady, 2008; Metsala, 1997), word production (Edwards, Beckman, & Munson, 2004; Munson, Swenson, & Manthei, 2005; Newman & German, 2005; Zamuner, Gerken, & Hammond, 2004), memory (Gathercole, Frankish, Pickering, & Peaker, 1999; Thomson, Richardson, & Goswami, 2005), and learning (Alt & Plante, 2006; Storkel, 2001, 2003, 2004a, 2009; Storkel, Armbrüster, & Hogan, 2006; Storkel & Maekawa, 2005; Swingley & Aslin, 2007). In a number of these studies, phonotactic probability and neighborhood density were calculated using readily available American English adult corpora and online calculators (Balota et al., 2007; Davis, 2005; Vitevitch & Luce, 2004), because comparable child calculators do not exist. However, the validity of the values generated from adult online calculators for child research warrants investigation. Moreover, an understanding of the relationship between values generated from child sources compared with those from adult sources is critical for developmental research, in which researchers seek to compare phonotactic probability and neighborhood density effects across different ages as the lexicon grows.

What evidence is there that child phonotactic probability and neighborhood density may differ from adult phonotactic probability and neighborhood density? To our knowledge, no researchers have investigated how phonotactic probability may change with development. How-

ever, numerous researchers have considered how neighborhood density may change from childhood to adulthood as the lexicon grows (Charles-Luce & Luce, 1990, 1995; Coady & Aslin, 2003; Dollaghan, 1994). Thus, we begin by examining what is known about neighborhood density changes and then apply the observed patterns to phonotactic probability. Across studies in which lexical growth was examined, there is clear evidence that the number of neighbors increases from childhood to adulthood, meaning that the child neighborhood density for a given word will tend to be lower than the adult neighborhood density for the same word (Charles-Luce & Luce, 1990, 1995; Coady & Aslin, 2003; Dollaghan, 1994). However, it is unknown whether these density differences are constant or variable across words or neighborhoods.

One possibility is that child neighborhood density differs from adult neighborhood density by a relatively constant value across neighborhoods. In this case, the difference in density for stimuli identified as sparse versus dense for one age group (e.g., children) will be approximately the same as that for an older age group (e.g., adults). Consider the following hypothetical example: The word *mouth*, with only 5 neighbors for children, is selected as a sparse word for children, and the word *tooth*, with 10 neighbors for children, is selected as a dense word for children. The difference between the sparse and dense condition is 5 neighbors for children. If growth in the lexicon is equally distributed across neighborhoods, two new neighbors may be added to each neighborhood, yielding 7 neighbors for adults for *mouth* and 12 neighbors for adults for *tooth*. Thus, although the child neighborhood density for each word (i.e., 5 and 10) is smaller than the adult

H. L. Storkel, hstorkel@ku.edu

neighborhood density for each word (i.e., 7 and 12), the difference between the sparse- and dense-neighborhood conditions is the same for both the child and the adult (i.e., the sparse neighborhood differs from the dense neighborhood by 5 neighbors).

In contrast, the difference between child and adult neighborhood density may not be the same across neighborhoods. In fact, there are good reasons to assume that this is the case. In particular, children and adults appear to learn dense-neighborhood words more readily than sparse-neighborhood words (Storkel, 2004a, 2009; Storkel et al., 2006), suggesting that more words may be added to dense neighborhoods than to sparse neighborhoods. To illustrate by continuing the previous hypothetical example, two new neighbors may be added to the neighborhood of *mouth*, yielding an adult neighborhood density of 7, whereas four new neighbors may be added to the neighborhood of *tooth*, yielding an adult neighborhood density of 14. Under this scenario, the difference between sparse- and dense-neighborhood stimuli is not the same for children (i.e., the sparse neighborhood differs from the dense neighborhood by five neighbors) and adults (i.e., the sparse neighborhood differs from the dense neighborhood by seven neighbors). Although the selected stimuli have sparse and dense neighborhoods for both age groups, it is unclear whether the neighborhoods are *equally* sparse and dense for each age group. This suggests the need to consider a measure of relative density when selecting stimuli for developmental research (cf. Coady & Aslin, 2003).

One relative measure of neighborhood density that may hold promise for addressing this issue is the *z* score. *z* scores express the distance from a reference point (i.e., the mean) in standard deviation units (i.e., [obtained value – *M*]/*SD*). They are commonly used in research to convert scores on different measurement scales to a common scale. In terms of the application to neighborhood density, *z* scores can quantify how extreme the neighborhood density of a given word is relative to those of other words in the child or adult lexicon and can express this in units that are comparable across lexicons of different sizes (i.e., standard deviation units). Thus, *z* scores provide a way of measuring whether the neighborhood of a given word is equally sparse or dense for each age group, using a common measurement scale. Storkel (2004b) described procedures for creating *z* scores for adult phonotactic probability and neighborhood density. Briefly, phonotactic probability and neighborhood density means and standard deviations were computed for all of the words of a given length in an adult corpus used for calculating phonotactic probability and neighborhood density, and these values were used to compute the *z* score for any given word. Thus, the *z* scores for *mouth* and *tooth* in our example might be –1.32 and –0.50, respectively. If we had a comparable reference point for children, we might find a *z* score of –1.24 for *mouth* and of –0.55 for *tooth* for children. In this scenario, the child *z* scores (i.e., –1.24 and –0.55) are relatively similar to the adult *z* scores (i.e., –1.32 and –0.50), even though the raw values appeared more discrepant. Therefore, in this example, we

might conclude that the word neighborhoods are equally sparse (or dense) for children and adults. Taken together, *z* scores hold promise for determining whether stimuli are equally sparse or dense for different age groups with different sized lexicons. However, it is unknown whether neighborhood density *z* scores would show the pattern in the hypothetical example, or whether discrepancies would be identified.

Although changes in phonotactic probability from childhood to adulthood have received less attention, the previously described scenarios for neighborhood density are relevant. As with neighborhood density, there is evidence that phonotactic probability influences word learning. In particular, when phonotactic probability is differentiated from neighborhood density, children and adults tend to learn rare sound sequences more readily than common sound sequences (Storkel, 2009; Storkel et al., 2006). The acquisition of rare sound sequences may lower phonotactic probability values for an adult relative to those of a child, and this lowering could be equivalently or asymmetrically distributed across rare and common sound sequences. It is also possible that differences in raw values could be eliminated when relative measures of phonotactic probability—such as *z* scores—are used.

Our goal in this research was to develop an online phonotactic probability and neighborhood density child calculator and to compare phonotactic probability and neighborhood density values based on this child calculator to those based on an adult calculator. The words used for comparing child and adult phonotactic probability and neighborhood density were the 380 nouns in the Words and Sentences Inventory of Fenson et al. (1993), because these words are likely to be known by both children and adults. The similarity between child and adult phonotactic probability and neighborhood density was examined via correlation to determine whether the relative ranking of the values was similar across children and adults. Child and adult phonotactic probability and neighborhood density also were compared via *t* test to determine whether the child values were significantly lower or higher than the adult values. On the basis of past work, child phonotactic probability was predicted to be significantly higher than adult phonotactic probability (Storkel, 2009; Storkel et al., 2006), and child neighborhood density was predicted to be significantly lower than adult neighborhood density (Charles-Luce & Luce, 1990, 1995; Coady & Aslin, 2003; Dollaghan, 1994). In the case of a significant *t* test, difference scores (i.e., adult value – child value) were computed for each significant variable and correlated with the child value for the same variable to determine whether the differences were equivalently distributed across the lexicon (i.e., no correlation between difference scores and child value) or not (i.e., significant correlation between difference scores and child value). On the basis of word-learning research, significant correlations between difference scores and child phonotactic probability (Storkel, 2009; Storkel et al., 2006) and child neighborhood density (Storkel, 2004a, 2009; Storkel et al., 2006) were expected. A similar set of analyses was performed for relative mea-

tures (i.e., z scores) of child and adult phonotactic probability and neighborhood density. Predictions were not made for this set of analyses, because few researchers have considered relative measures of phonotactic probability or neighborhood density (but see Coady & Aslin, 2003), and none have used z scores.

METHOD

Child Corpus

Words and frequency of occurrence were taken from two spoken language corpora: Kolson (1960), based on words produced by kindergarten children at home, school, and in an elicited picture task ($n = 3,728$ different words), and Moe, Hopkins, and Rush (1982), based on words produced by first-grade children during an examiner-led interview ($n = 6,412$ different words). Hard-copy data from the original publications were scanned into a computer and converted to text. The scanned data were checked for accuracy by comparing the electronic data to the hard-copy data. The two databases were then combined. Words appearing in both databases were combined into one entry by adding the frequency from each database. The remaining procedures were undertaken so that the child corpus would better match the existing adult corpus, which is based on a dictionary. First, the $\log_{10}$ of the combined raw frequency counts was calculated. A constant of 1.0 was added to the log value to avoid values of 0, which occur when the raw frequency is 1. Second, a computer-readable phonemic transcription (i.e., pronunciation) of each word was obtained from the existing adult corpus. If a word was not available in the adult corpus, a dictionary was consulted (Longman, 1983). Dictionary transcriptions were altered to follow transcription conventions in the adult corpus, which tended to make greater use of syllabic sonorant consonants. In the event that a word was not found in the dictionary, a native speaker of English provided a transcription for the word following the transcription conventions in the adult corpus. This occurred primarily for proper names (e.g., *Allison*, *Frankenstein*, *Duluth*). Finally, similar forms of words were eliminated for two types of words: words that were grammatically related (i.e., inflected vs. uninflected forms) and words that are pronounced the same but differ in meaning and spelling (i.e., homophones). These procedures are in line with current theories of the lexicon, which assume that only one form of a word is stored and which serve to make the child corpus more comparable to the available adult corpus (described below). More specifically, all inflected forms of a word (e.g., *running*) were eliminated from the corpus, leaving only the uninflected form (e.g., *run*). This included consideration of both grammatical and ungrammatical inflected forms (e.g., *deers*). The only exception to this procedure was that inflected forms were retained in the corpus if the uninflected form was absent (e.g., *aces* was retained because *ace* was absent). Likewise, homophonous word forms (e.g., *by*, *bye*, *buy*) were collapsed into one form. These procedures yielded a corpus of 4,832 different words (i.e., types) and 1,028,417 total words (i.e., tokens).

Adult Corpus

The adult corpus for comparison was the Hoosier Mental Lexicon (Nusbaum, Pisoni, & Davis, 1984). This corpus consists of 19,290 different words (i.e., types) from a dictionary (*The New Merriam-Webster Pocket Dictionary*, 1964). For the present study, only the computer-readable phonemic transcription and written word frequency (Kučera & Francis, 1967) were used. In terms of overlap between the two corpora, 58% of the words in the child corpus appeared in the adult corpus. In addition, words in the child corpus ($M = 4.87$ sounds, $SD = 1.77$ sounds, range = 1–14 sounds) were significantly shorter than the words in the adult corpus ($M = 6.35$ sounds, $SD = 2.31$ sounds, range = 1–15 sounds) [$t(24120) = -41.55, p < .001$]. Words in the child corpus ($M = 2.04, SD = 0.87$,

range = 1.00–5.73) also were significantly higher in log frequency than the words in the adult corpus ($M = 1.49, SD = 0.69$, range = 1.00–5.84) [$t(24120) = 47.71, p < .001$].

Online Calculator

An online interface was created to calculate phonotactic probability and neighborhood density using either the child or the adult corpora. Although online calculators for adult corpora already exist, the adult corpus was included in this calculator so that the same interface for computing phonotactic probability and neighborhood density could be applied to each corpus. That is, any programming errors would affect both the child and the adult values. The calculator is available at www.bncdnet.ku.edu/cml/info_ccc.vi. The prototype software for the calculator was built in LabVIEW as a stand-alone application with a graphical interface and is compatible with Linux, Mac, and Windows operating systems. The Web version uses this same application with the LabVIEW Internet Toolkit, CGI, and an Apache 2 HTTP server on a Linux system. The Web interface has been tested with versions of Mozilla Firefox and Microsoft Internet Explorer browsers and requires no additional browser plug-ins.

To use the calculator, a phonemic transcription of a target word (or words) in a computer-readable format is entered in a textbox with one word per line, the child or adult corpus is selected, and output variables are selected (described below). Additional detailed instructions and an example item are provided on the calculator Web site.

Phonotactic probability algorithm. Two measures of phonotactic probability were computed for this analysis: positional segment average and biphone average (Storkel, 2004b; Vitevitch & Luce, 2004). In addition to these two summary variables, the calculator will return the component values that lead to these summary variables, as is detailed below. The calculation for positional segment average begins by computing the positional segment frequency, referred to as *Pos Seg Freq* in the calculator interface. The *positional segment frequency* is computed for each sound in the target word by iterating over every entry in the corpus that is long enough to have any sound in the corresponding position (counted from the left edge of the word without respect to syllable structure) and by checking for matches against each sound in the target word. The log frequency of all of the words in the corpus that contain the given sound in a given word position (identified as S1, S2, etc.) is summed and then divided by the sum of the log frequency of all of the words in the corpus that contain any sound in the given word position. Thus, the number of positional segment frequencies returned will correspond to the number of sounds in the target word. The *positional segment sum* (i.e., *Pos Sum* in the calculator interface) is computed by adding the positional segment frequency for each sound in the target word. The *positional segment average* (i.e., *Pos Ave* in the calculator interface) is computed by dividing the positional segment sum by the number of sounds in the word.

The biphone average is computed in a similar manner, except that pairs of adjacent sounds are used in the calculations, rather than individual sounds. The calculation for biphone average begins by computing the *biphone frequency* (i.e., *Biphone Freq* in the calculator interface), which is computed for each pair of sounds in the target word by summing the log frequency of all of the words in the corpus that contain the given pair of sounds in a given word position (identified as B1, B2, etc.) and then dividing by the sum of the log frequency of all of the words in the corpus that contain any sound in the given word position. Thus, the number of biphone frequencies returned will correspond to the number of sound pairs in the target word. The *biphone sum* (i.e., *Biphone Sum* in the calculator interface) is computed by adding the biphone frequency for each sound pair in the target word. The *biphone average* (i.e., *Biphone Ave* in the calculator interface) is computed by dividing the biphone sum by the number of sound pairs in the word (which will be the number of sounds minus 1).

Neighborhood density algorithm. Neighborhood density, referred to as *Num of Nbrs* in the calculator interface, is calculated

by counting all of the words in the corpus that differ from the target word by a one sound substitution, addition, or deletion in any word position (Balota et al., 2007; Storkel, 2004b). To find neighbors, the calculator creates a regular expression and then finds every matching entry in the corpus. For example, to find an addition in the second position for /sɪŋ/, which would be entered as *sIG* in the calculator input box, the regular expression *s.IG* is formed with the wild card character (.) in the second position. The wild card character indicates that only one sound may be inserted in the second position, so /sɪŋ/ would be identified as a match but /striŋ/ would not be, because two sounds are inserted in the second position. This process is iterative such that multiple regular expressions are created by the calculator program to search for all possible substitutions, deletions, and additions in any word position. The calculator will show the neighbors by checking the Show Neighbors box. The calculator also provides other details of neighborhood structure (see the details at the Web site).

Phonotactic probability and neighborhood density characteristics of the child and adult corpora are summarized in Appendix A.

Words for Analysis

The words used for comparing child and adult phonotactic probability and neighborhood density were the 380 nouns in the Words and Sentences Inventory of Fenson et al. (1993), because these words are likely known by both children and adults and have been used in past word-learning research (Storkel, 2004a, 2009). In addition, the words in this set vary in the sounds targeted (i.e., all English sounds present), syllable structure (e.g., CV, CCV, CCCV, VC, VCC, CVC, CCVC, CCCVC, CVCC), and word length (i.e., $M = 4.42$ sounds, $SD = 1.56$ sounds, range = 1–10 sounds), suggesting that the set represents a wide range of word structures learned by young children.

For each of the 380 nouns, positional segment average, biphone average, and neighborhood density were computed using the online calculator and the child corpus. In addition to these raw values, z scores were computed for each word, following the procedures of Storkel (2004b; [obtained value – M]/ SD) and using the child means and standard deviations for the appropriate word length (see Appendix A). The same raw values were computed for each word using the online calculator and the adult corpus, and then z scores were calculated using the adult means and standard deviations for the appropriate word length (see Appendix A).

A parallel analysis was completed for a set of 310 nonwords; the results are shown in Appendix B. Generally, the nonword analysis produced results similar to those of the real-word analysis.

RESULTS AND DISCUSSION

Raw Values

The first issue to be addressed by this study was the relationship between child and adult phonotactic probability and neighborhood density raw values. The results show that child and adult raw values were significantly positively correlated for positional segment averages [$r(380) = .90$, $p < .001$], biphone average [$r(378) = .89$, $p < .001$], and neighborhood density [$r(380) = .94$, $p < .001$]. As is shown in Figure 1, the child and adult values for each variable tended to decrease or increase in tandem.

Turning to t -test comparisons, the child positional segment averages ($M = 0.051$, $SD = 0.015$, range = 0.006–0.096) were significantly higher than the adult positional segment averages ($M = 0.047$, $SD = 0.014$, range = 0.004–0.089) [$t(379) = 12.47$, $p < .001$, $\eta_p^2 =$

Figure 1. Scatterplots of child versus adult positional segment average (top), biphone average (middle), and neighborhood density (bottom). The solid line indicates the linear regression fit line. The dashed line is a reference line indicating a perfect correlation.

.29]. Likewise, the child biphone averages ($M = 0.0044$, $SD = 0.0027$, range = 0.0002–0.0146) were significantly higher than the adult biphone averages ($M = 0.0038$, $SD = 0.0026$, range = 0.0001–0.0164) [$t(377) = 10.56$, $p < .001$, $\eta_p^2 = .23$]. These findings are consistent with a priori predictions that ease of learning of rare sound sequences, as has been reported in previous research (Storkel, 2009; Storkel et al., 2006), would lead to a lowering of phonotactic probability from childhood to adulthood. Another factor that accounts for the difference between child and adult phonotactic probability is that the words in the child corpus were significantly higher in frequency than the words in the adult corpus. Because word frequency is used in phonotactic probability computations, the addition of low-frequency words to the lexicon would also lead to a lowering of phonotactic probability from childhood to adulthood.

The opposite pattern was observed for neighborhood density. Here, child neighborhood density ($M = 5.5$, $SD = 6.6$, range = 0–34) was significantly lower than adult neighborhood density ($M = 8.7$, $SD = 9.2$, range = 0–40) [$t(379) = -16.59$, $p < .001$, $\eta_p^2 = .42$]. This finding is consistent with past studies showing that child neighborhood densities are significantly lower than adult neighborhood densities because of differences in the overall size of the lexicon of children versus that of adults (Charles-Luce & Luce, 1990, 1995; Coady & Aslin, 2003; Dollaghan, 1994).

Difference Scores

The second issue to be addressed was whether the observed significant differences between child and adult phonotactic probability, as well as those between child and adult neighborhood density, were equivalently distributed across the lexicon. To address this issue, the relationship between these child and adult differences and each child variable was explored. Difference scores were computed for each variable by subtracting the child value from the adult value. In this way, negative difference scores indicate that the child value is larger than the adult value, and positive difference scores indicate that the child value is smaller than the adult value.

For phonotactic probability, difference scores were significantly negatively correlated with child positional segment average [$r(380) = -.35$, $p < .001$] and child biphone average [$r(378) = -.33$, $p < .001$]. As is shown in Figure 2, the difference scores decreased (i.e., became more negative) as child positional segment or biphone averages increased. That is, differences between child phonotactic probability and adult phonotactic probability were not equivalently distributed across the lexicon. Instead, the differences between child and adult positional segment averages were smaller for lower probability sound sequences (i.e., those at or below the child mean; $M_{\text{difference}} = -0.0022$, $SD = 0.0064$, range = -0.0016 – 0.0288) than for higher probability sound sequences (i.e., those above the child mean; $M_{\text{difference}} = -0.0062$, $SD = 0.0062$, range = -0.0250 – 0.0069). Likewise, differences

Figure 2. Scatterplots of difference scores (adult – child) versus child positional segment average (top), biphone average (middle), and neighborhood density (bottom). The solid line indicates the linear regression fit line. The dashed line is a reference line indicating a difference score of zero (i.e., adult value = child value). Points falling below the line (i.e., a negative difference score) indicate that the child value is higher than the adult value. Points falling above the line (i.e., a positive difference score) indicate that the child value is lower than the adult value.

between child and adult biphone averages were smaller for lower probability sound sequences ($M_{\text{difference}} = -0.0004$, $SD = 0.0010$, range = -0.0027 – 0.0046) than for higher probability sound sequences ($M_{\text{difference}} = -0.0010$, $SD = 0.0014$, range = -0.0055 – 0.0027). This suggests that the learning of rare sound sequences (Storkel, 2009; Storkel et al., 2006) and low-frequency words that occurred with development had a greater impact on higher probability sound sequences than on lower probability sound sequences.

For neighborhood density, difference scores were significantly positively correlated with child neighborhood density [$r(380) = .54$, $p < .001$]. As is shown in Figure 2, difference scores increased (i.e., became more positive) as neighborhood density increased. As with phonotactic probability, the differences between child and adult neighborhood density were not equivalently distributed across the lexicon. Rather, the differences between child and adult densities were smaller for sparser neighborhoods (i.e., those at or below the child mean; $M = 1.37$ neighbors, $SD = 2.22$ neighbors, range = -3 – 11) than for denser neighborhoods (i.e., those above the child mean; $M = 5.20$ neighbors, $SD = 4.07$, range = -4 – 17). As was predicted from word-learning research (Storkel, 2004a, 2009; Storkel et al., 2006), more words were added to dense neighborhoods than to sparse neighborhoods as the lexicon grew, leading to larger discrepancies between the child and adult neighborhood densities for dense neighborhoods than for sparse neighborhoods.

z Scores

The third issue to be addressed was whether the patterns observed with raw values would also be observed for relative measures of phonotactic probability and neighborhood density (i.e., z scores). The same set of analyses performed for the raw values was performed for the z-score data. As with the raw values, the child and adult z scores were significantly positively correlated for positional segment average z scores [$r(380) = .88$, $p < .001$], biphone average z scores [$r(378) = .87$, $p < .001$], and neighborhood density z scores [$r(380) = .71$, $p < .001$]. As is shown in Figure 3, the child and adult z scores tended to increase or decrease in tandem.

Turning to *t*-test analyses, the child positional segment average z scores ($M = 0.02$, $SD = 0.95$, range = -2.64 – 2.89) were significantly lower than the adult positional segment average z scores ($M = 0.09$, $SD = 0.97$, range = -2.21 – 2.95) [$t(379) = -2.75$, $p < .01$, $\eta_p^2 = .02$] (refer to Figure 3). In contrast, the child biphone average z scores ($M = 0.03$, $SD = 1.03$, range = -1.89 – 4.11) were significantly higher than the adult biphone average z scores ($M = -0.05$, $SD = 0.99$, range = -1.82 – 5.20) [$t(377) = 2.98$, $p < .01$, $\eta_p^2 = .02$] (refer to Figure 3). Although these differences are statistically significant, note that the effect sizes are small and that the difference scores were relatively small for both positional segment averages ($M = 0.07$, $SD = 0.48$, range = -2.69 – 2.00) and biphone averages ($M = -0.08$, $SD = 0.52$, range = -1.99 – 2.16). Thus, the distributions of child and adult z scores largely overlapped.

Figure 3. Scatterplots of child versus adult positional segment average z scores (top), biphone average z scores (middle), and neighborhood density z scores (bottom). The solid line indicates the linear regression fit line. The dashed line is a reference line indicating a perfect correlation.

Figure 4. Scatterplots of z-score difference (adult – child) relative to the child positional segment average z score (top), the child biphone average z score (middle), and the child neighborhood density z score (bottom). The solid line indicates the linear regression fit line. The horizontal dashed line is a reference line indicating a difference score of zero (i.e., adult value = child value). Points falling below the line (i.e., a negative difference score) indicate that the child value is higher than the adult value. Points falling above the line (i.e., a positive difference score) indicate that the child value is lower than the adult value. The vertical dashed line indicates a z score of 0.00, differentiating low probability or density (values below 0.00) from high (values greater than 0.00).

Moreover, an analysis of z-score differences (i.e., adult z score – child z score) showed a pattern that resembled regression to the mean. The z-score differences (i.e., adult z score – child z score) were significantly negatively correlated with child positional segment average z scores [$r(380) = -.22, p < .001$] and with child biphone average z scores [$r(378) = -.33, p < .001$]. Figure 4 shows difference scores relative to child z scores. When the child z scores were negative (i.e., low probability), the adult z scores tended to be less negative for positional segment average ($M_{\text{difference}} = 0.14, SD = 0.43, \text{range} = -1.42\text{--}1.31$) and for biphone average ($M_{\text{difference}} = 0.06, SD = 0.43, \text{range} = -1.29\text{--}2.16$). In contrast, when the child z scores were positive (i.e., high probability), the adult z scores tended to be less positive for positional segment average ($M_{\text{difference}} = -0.01, SD = 0.51, \text{range} = -2.69\text{--}4.69$) and for biphone average ($M_{\text{difference}} = -0.27, SD = 0.56, \text{range} = -1.99\text{--}1.35$). In both cases, the adult z scores tended to be less extreme than the child z scores, a pattern that is consistent with regression to the mean rather than developmental differences between children and adults. Note that the concept of regression to the mean is typically applied to multiple observations or measures on a single person, with the underlying hypothesis being that each observation reflects the person's true score plus some amount of error that will variably influence the obtained score on repeated administration. The effect of the variability is that extreme scores on the first administration will tend to become less extreme on repeated administration. Although the present case is a bit different (i.e., multiple measures on a single set of words), the logic is similar. That is, there is a true phonotactic probability score for each word, but there is some amount of error associated with each corpus (or method) used to calculate phonotactic probability. The error across corpora will variably influence the obtained phonotactic probability for a given word such that extreme scores based on one corpus will tend to become less extreme in a second corpus.

Turning to neighborhood density z scores, the child z scores ($M = 0.04, SD = 1.04, \text{range} = -2.00\text{--}4.73$) did not differ significantly from the adult z scores ($M = 0.05, SD = 0.96, \text{range} = -2.03\text{--}4.20$) [$t(379) = -0.336, p > .70, \eta_p^2 < .001$]. Likewise, z-score differences (i.e., adult z score – child z score) were relatively small ($M = 0.01, SD = 0.76, \text{range} = -5.08\text{--}2.29$). As is shown in Figure 4, differences between the child and adult neighborhood density z scores were consistent with a pattern of regression to the mean.

SUMMARY AND CONCLUSION

Child and adult phonotactic probability and neighborhood density were significantly correlated for both raw values and z scores. This suggests that the corpora used to compute phonotactic probability and neighborhood density may not be critical when only gross or extreme distinctions are made (e.g., low vs. high probability or density). In contrast, a *t*-test analysis of raw values showed significant discrepancies between the child and adult values. Specifi-

cally, the child phonotactic probability was higher than the adult phonotactic probability, whereas the child neighborhood density was lower than the adult neighborhood density. Moreover, the differences between the child and adult phonotactic probability and neighborhood density were not distributed equivalently across the lexicon. In particular, the child and adult differences were larger for higher probability sound sequences and for denser neighborhoods. Thus, when more precise or fine-grained distinctions need to be made, it may be more critical to select corpora that are more representative of the words known by the study population. In addition, *z* scores may be useful in establishing whether the words are equivalently rare or common, sparse or dense, given the hypothesized size of the lexicon of the study population. Use of *z* scores appeared to reduce the differences between the child and adult values for phonotactic probability and neighborhood density. Moreover, *z* scores would be useful in comparing phonotactic probability and neighborhood density across studies in which different corpora are used to determine whether stimuli are equivalently rare or common, sparse or dense, across studies. The results from the analysis of nonwords (see Appendix B) generally parallel these results from real words, bolstering these conclusions.

AUTHOR NOTE

This research was supported by NIH Grants DC08095, DC00052, DC009135, DC05803, and HD02528. Douglas S. Kieweg created the Web interface and program for the online calculator. The following individuals contributed to development of the child corpus and/or data processing and analysis: Shinying Chu, Maki Sueto, and Ashlee Widler. J.R.H. is now at Indiana University, Bloomington. Correspondence concerning this article should be addressed to H. L. Storkel, Department of Speech-Language-Hearing: Sciences and Disorders, University of Kansas, 3001 Dole Human Development Center, 1000 Sunnyside Avenue, Lawrence, KS 66045-7555 (e-mail: hstorkel@ku.edu).

REFERENCES

- ALT, M., & PLANTE, E. (2006). Factors that influence lexical and semantic fast mapping of young children with specific language impairment. *Journal of Speech, Language, & Hearing Research*, **49**, 941-954.
- BALOTA, D. A., YAP, M. J., CORTESE, M. J., HUTCHISON, K. A., KESSLER, B., LOFTIS, B., ET AL. (2007). The English Lexicon Project. *Behavior Research Methods*, **39**, 445-459.
- CHARLES-LUCE, J., & LUCE, P. A. (1990). Similarity neighborhoods of words in young children's lexicons. *Journal of Child Language*, **17**, 205-215.
- CHARLES-LUCE, J., & LUCE, P. A. (1995). An examination of similarity neighborhoods in young children's receptive vocabularies. *Journal of Child Language*, **22**, 727-735.
- COADY, J. A., & ASLIN, R. N. (2003). Phonological neighborhoods in the developing lexicon. *Journal of Child Language*, **30**, 441-469.
- DAVIS, C. J. (2005). *N-Watch*: A program for deriving neighborhood size and other psycholinguistic statistics. *Behavior Research Methods*, **37**, 65-70.
- DOLLAGHAN, C. A. (1994). Children's phonological neighborhoods: Half empty or half full? *Journal of Child Language*, **21**, 257-271.
- EDWARDS, J., BECKMAN, M., & MUNSON, B. (2004). The interaction between vocabulary size and phonotactic probability effects on children's production accuracy and fluency in nonword repetition. *Journal of Speech, Language, & Hearing Research*, **47**, 421-436.
- FENSON, L., DALE, P. S., REZNICK, J. S., THAL, D., BATES, E., HARTUNG, J. P., ET AL. (1993). *The MacArthur Communicative Development Inventories: User's guide and technical manual*. San Diego: Singular.
- GARLOCK, V. M., WALLEY, A. C., & METSALA, J. L. (2001). Age-of-acquisition, word frequency, and neighborhood density effects on spoken word recognition by children and adults. *Journal of Memory & Language*, **45**, 468-492.
- GATHERCOLE, S. E., FRANKISH, C. R., PICKERING, S. J., & PEAKER, S. (1999). Phonotactic influences on short-term memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**, 84-95.
- KOLSON, C. J. (1960). *The vocabulary of kindergarten children*. Unpublished doctoral dissertation, University of Pittsburgh.
- KUČERA, H., & FRANCIS, W. N. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.
- Longman dictionary of American English* (1983). White Plains, NY: Longman.
- MAINELA-ARNOLD, E., EVANS, J. L., & COADY, J. A. (2008). Lexical representations in children with SLI: Evidence from a frequency-manipulated gating task. *Journal of Speech, Language, & Hearing Research*, **51**, 381-393.
- METSALA, J. L. (1997). An examination of word frequency and neighborhood density in the development of spoken-word recognition. *Memory & Cognition*, **25**, 47-56.
- MOE, A. J., HOPKINS, K. J., & RUSH, R. T. (1982). *The vocabulary of first grade children*. Springfield, IL: Thomas.
- MUNSON, B., SWENSON, C. L., & MANTHEI, S. C. (2005). Lexical and phonological organization in children: Evidence from repetition tasks. *Journal of Speech, Language, & Hearing Research*, **48**, 108-124.
- The new Merriam-Webster pocket dictionary* (1964). New York: Pocket.
- NEWMAN, R. S., & GERMAN, D. J. (2005). Life span effects of lexical factors on oral naming. *Language & Speech*, **48**, 123-156.
- NUSBAUM, H. C., PISONI, D. B., & DAVIS, C. K. (1984). *Sizing up the Hoosier mental lexicon* (Research on Spoken Language Processing Rep. No. 10, pp. 357-376). Bloomington, IN: Speech Research Laboratory, Indiana University.
- STORKEL, H. L. (2001). Learning new words: Phonotactic probability in language development. *Journal of Speech, Language, & Hearing Research*, **44**, 1321-1337.
- STORKEL, H. L. (2003). Learning new words II: Phonotactic probability in verb learning. *Journal of Speech, Language, & Hearing Research*, **46**, 1312-1323.
- STORKEL, H. L. (2004a). Do children acquire dense neighborhoods? An investigation of similarity neighborhoods in lexical acquisition. *Applied Psycholinguistics*, **25**, 201-221.
- STORKEL, H. L. (2004b). Methods for minimizing the confounding effects of word length in the analysis of phonotactic probability and neighborhood density. *Journal of Speech, Language, & Hearing Research*, **47**, 1454-1468.
- 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., ARMBRÜSTER, J., & HOGAN, T. P. (2006). Differentiating phonotactic probability and neighborhood density in adult word learning. *Journal of Speech, Language, & Hearing Research*, **49**, 1175-1192.
- STORKEL, H. L., & MAEKAWA, J. (2005). A comparison of homonym and novel word learning: The role of phonotactic probability and word frequency. *Journal of Child Language*, **32**, 827-853.
- SWINGLEY, D., & ASLIN, R. N. (2007). Lexical competition in young children's word learning. *Cognitive Psychology*, **54**, 99-132.
- THOMSON, J. M., RICHARDSON, U., & GOSWAMI, U. (2005). Phonological similarity neighborhoods and children's short-term memory: Typical development and dyslexia. *Memory & Cognition*, **33**, 1210-1219.
- VITEVITCH, M. S., & LUCE, P. A. (2004). A Web-based interface to calculate phonotactic probability for words and nonwords in English. *Behavior Research Methods, Instruments, & Computers*, **36**, 481-487.
- ZAMUNER, T. S., GERKEN, L., & HAMMOND, M. (2004). Phonotactic probabilities in young children's speech production. *Journal of Child Language*, **31**, 515-536.

APPENDIX A

Phonotactic Probability and Neighborhood Density As a Function of Word Length for Child and Adult Corpora

Word Length in Phonemes	Number of Words per Corpus		Positional Segment Average				Biphone Average				Neighborhood Density			
			Child		Adult		Child		Adult		Child		Adult	
	Child	Adult	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
1	5	5	.0082	.0042	.0113	.0118	—	—	—	—	16.80	7.56	21.40	8.62
2	157	194	.0335	.0201	.0305	.0170	.0018	.0014	.0015	.0028	17.29*	7.18	22.89*	8.05
3	958	1,608	.0512*	.0175	.0449*	.0166	.0040*	.0026	.0030*	.0026	12.31*	6.14	18.21*	8.47
4	1,286	2,899	.0545*	.0147	.0466*	.0135	.0049*	.0027	.0040*	.0027	4.06*	3.04	6.87*	4.53
5	943	3,114	.0515*	.0139	.0462*	.0121	.0046*	.0025	.0042*	.0025	0.91*	1.23	1.78*	1.96
6	639	3,039	.0498*	.0126	.0474*	.0118	.0043	.0022	.0046	.0025	0.22*	0.50	0.53*	0.86
7	442	2,725	.0486	.0112	.0483	.0111	.0040*	.0018	.0049*	.0023	0.09*	0.30	0.25*	0.54
8	210	2,202	.0493	.0104	.0504	.0108	.0041*	.0017	.0057*	.0026	0.05*	0.24	0.16*	0.43
9	117	1,534	.0512	.0101	.0521	.0104	.0049*	.0018	.0069*	.0036	0.04	0.20	0.15	0.43
10	44	1,019	.0525	.0107	.0570	.0114	.0052*	.0018	.0093*	.0059	0.09	0.29	0.12	0.34
11	23	526	.0537	.0145	.0564	.0111	.0055*	.0021	.0082*	.0043	0.09	0.29	0.06	0.24
12	4	266	.0708	.0103	.0603	.0118	.0101	.0015	.0094	.0054	0.00	0.00	0.03	0.17
13	2	106	.0743	.0330	.0563	.0108	.0182*	.0066	.0077*	.0036	0.00	0.00	0.00	0.00
14	2	37	.1066	.0416	.0563	.0113	.0324	.0135	.0080	.0034	0.00	0.00	0.00	0.00
15	0	16	—	—	.0629	.0114	—	—	.0099	.0024	—	—	0.00	0.00
All words	4,832	19,290	.0511*	.0151	.0485*	.0129	.0044*	.0026	.0051*	.0034	4.32*	6.16	3.23*	6.39

Note—Although the values were recomputed for this study, adult values match those reported in Table 2 of Storkel (2004b), *Journal of Speech, Language, and Hearing Research*, 47, p. 1460. Copyright 2004 by the American Speech-Language-Hearing Association. Reprinted with permission. * $p < .001$.

APPENDIX B

Analysis of a Set of Nonwords

Nonwords for Analysis

The nonwords used for comparing child and adult phonotactic probability/neighborhood density were 310 nonwords developed in our lab for published, unpublished, and ongoing research studies on word learning and word representations. Of these 310 nonwords, 298 were one-syllable consonant-vowel-consonant (CVC) sequences, and 12 were two-syllable CVCVC sequences. Note that these stimuli do not represent a random sample of CVC nonwords, because various constraints (e.g., consonants needed to be acquired early) were imposed during stimuli creation, depending on the goals of the study, and, in most cases, the stimuli were selected for extreme values of phonotactic probability and/or neighborhood density based on the adult corpus (i.e., low vs. high probability and/or density was an independent variable in the study).

Raw Values

The results of this set of analyses parallel the findings reported for real words: child and adult values were significantly correlated (see Table B1); child phonotactic probability was significantly higher than adult; child neighborhood density was significantly lower than adult (see Table B2).

Table B1
Correlation Between Child and Adult Values

	<i>r</i>	<i>p</i>
Positional segment average	.92	<.001
Biphone average	.85	<.001
Neighborhood density	.90	<.001

Table B2
t Test Between Child and Adult Values

	<i>t</i>	<i>p</i>	η_p^2	Child			Adult		
				<i>M</i>	<i>SD</i>	Range	<i>M</i>	<i>SD</i>	Range
Positional segment average	15.74	<.001	.45	0.047	0.012	0.015–0.082	0.042	0.012	0.014–0.085
Biphone average	6.99	<.001	.14	0.0026	0.0017	0.0000–0.0096	0.0022	0.0018	0.0001–0.0133
Neighborhood density	–28.97	<.001	.73	9.6	5.3	0–26	15.33	7.4	0–39

(Continued on next page)

APPENDIX B (Continued)

Difference Scores

The results of this set of analyses parallel the findings reported for real words: The child and adult differences were larger for higher probability sequences and for denser neighborhoods (see Table B3).

***z* Scores**

The results of the correlation analysis parallel the findings reported for real words: The child and adult *z* scores were significantly correlated (see Table B4). In contrast, results of the *t*-test analysis differed from the findings reported for real words (see Table B5). Specifically, significant differences were obtained for all three variables, with the child *z* scores being significantly lower than the adult *z* scores for each variable. However, similar to the analysis of real words, the effect sizes were lower for *z* scores than for raw values, and the differences between the child and adult values were relatively small for *z* scores. Also similar to the analysis of real words, when the *z*-score differences were significantly correlated with the child *z* scores (see Table B6), the pattern was consistent with an interpretation of regression to the mean.

Table B3
Correlation of Difference Scores (Adult – Child) With Each Child Variable

	<i>r</i>	<i>p</i>	Low			High		
			<i>M</i>	<i>SD</i>	Range	<i>M</i>	<i>SD</i>	Range
Positional segment average	–.12	.04	– 0.004	0.005	–0.018 to 0.009	– 0.005	0.005	–0.019 to 0.007
Biphone average	–.13	.03	– 0.0004	0.0007	–0.0021 to 0.0042	– 0.0004	0.0018	–0.0048 to 0.0040
Neighborhood density	.39	<.001	5.2	3.1	–1 to 15	7.1	3.9	–1 to 16

Table B4
Correlation Between Child and Adult *z* Scores

	<i>r</i>	<i>p</i>
Positional segment average	.92	<.001
Biphone average	.86	<.001
Neighborhood density	.88	<.001

Table B5
***t* Test Between Child and Adult *z* Scores**

	<i>t</i>	<i>p</i>	η_p^2	Child			Adult		
				<i>M</i>	<i>SD</i>	Range	<i>M</i>	<i>SD</i>	Range
Positional segment average	–5.94	<.001	.10	– 0.26	0.71	–2.65 to 1.77	– 0.16	0.78	–2.69 to 2.42
Biphone average	–10.95	<.001	.28	– 0.55	0.65	–1.76 to 2.15	– 0.32	0.71	–1.64 to 3.96
Neighborhood density	–4.43	<.001	.06	– 0.39	0.80	–1.84 to 2.23	– 0.29	0.80	–2.03 to 2.45

Table B6
Correlation of *z*-Score Differences (Adult – Child) With Each Child Variable

	<i>r</i>	<i>p</i>	All			Low			High		
			<i>M</i>	<i>SD</i>	Range	<i>M</i>	<i>SD</i>	Range	<i>M</i>	<i>SD</i>	Range
Positional segment average	.03	.62	0.10	0.30	–0.73 to 1.08	0.10	0.29	–0.73 to 1.08	0.11	0.32	–0.73 to 0.84
Biphone average	–.11	.05	0.23	0.37	–1.46 to 2.00	0.23	0.28	–0.42 to 2.00	0.22	0.71	–1.46 to 1.92
Neighborhood density	–.26	<.001	0.10	0.39	–0.94 to 1.19	0.15	0.34	–0.58 to 1.19	– 0.05	0.46	–0.94 to 0.93

(Manuscript received October 31, 2009;
revision accepted for publication January 18, 2010.)