

Acquiring an artificial lexicon: Segment type and order information in early lexical entries

Sarah C. Creel *, Richard N. Aslin, Michael K. Tanenhaus

Department of Brain and Cognitive Sciences, Meliora Hall, University of Rochester, Rochester, NY 14627, USA

Received 20 December 2004; revision received 19 September 2005

Available online 28 November 2005

Abstract

The role of segment similarity in early (i.e., partially learned) lexical entries was assessed using artificial lexicons in a referential context. During a learning phase participants heard 40 nonsense words, each accompanied by an unfamiliar picture. In testing, participants heard the direction “Click on the [X]”, and chose which of four pictures was the target (X). Target lexical items (e.g., *pibo*) appeared with foils that were similar: cohort items (*pibu*), rhymes (*dibo*), matched consonants (*pabu*) or matched vowels (*diko*). Two initial experiments demonstrated cohort and rhyme confusions, similar to lexical activation findings. Four further experiments explored the role of segment similarity in word confusions. Consonant-matched CVCV stimuli were more strongly confused with each other than were vowel-matched CVCV stimuli. Placing consonants in syllable-final position (VC[f]VC) weakened consonant effects and strengthened vowel effects. These results suggest that syllable-initial segments play a strong role in word similarity and constrain the organization of new lexical items.

© 2005 Elsevier Inc. All rights reserved.

Keywords: Artificial lexicon; Lexical learning; Lexical acquisition; Consonant; Vowel; Word learning; Lexical activation

The research reported here addresses two issues about the representation and processing of spoken words. The first issue is whether vowels and consonants contribute differentially to lexical similarity and if so, whether differences are tied to the potential information provided by each class of segment and/or the position they typically occupy in syllables. In order to address this issue, we examine confusions between newly learned words in an artificial lexicon. Use of an artificial lexicon allows us to manipulate properties that are otherwise confounded in natural language lexicons. For example,

most languages have more consonants than vowels and consonants are more common than vowels in word and syllable-initial position.

The second issue is whether there are qualitative differences in the relative importance of consonants and vowels early in learning compared to when a word has been well learned. Partially learned lexical representations, which we will refer to as “early” lexical entries, may show the onset-based processing biases characteristic of words in natural language lexicons (Marslen-Wilson, 1987) and in artificial language lexicons (Magnuson, Tanenhaus, Aslin, & Dahan, 2003) or they may be represented more holistically. If there are qualitative changes in representations as words become better learned, then the relative importance of vowels and consonants might change with the degree of learning.

* Corresponding author. Present address: Psychology Department, University of Pennsylvania, Philadelphia, PA 19104, USA. Fax: +1 215 898 7301.

E-mail address: creel@psych.upenn (S.C. Creel).

In the remainder of the introduction, we first motivate the rationale for examining differences among consonants and vowels and then introduce the rationale for using the artificial lexicon methodology. We then consider the issue of whether early lexical entries are represented holistically or in the more phonologically elaborated, onset-biased way that characterizes well-learned natural lexical representations.

Research on spoken word recognition has shown that as the acoustic properties of words unfold in time, multiple word candidates are activated in parallel and compete for recognition. Early in processing, the most active candidates have similar onsets (Marslen-Wilson, 1987; Marslen-Wilson & Zwitserlood, 1989; Marslen-Wilson, Moss, & van Halen, 1996), with activation modulated by frequency (Dahan, Magnuson, & Tanenhaus, 2001), among other factors. However, as more input arrives, words that mismatch at onset but are globally similar to the target word also become activated (Alloppenna, Magnuson, & Tanenhaus, 1998; Connine, Titone, & Wang, 1993; Goldinger, 1998; Luce & Pisoni, 1998), though these words never compete as strongly as those with similar onsets.

The lexical candidates that become most active when a spoken word is processed define the competitor or processing environment for that word. Although effects of acoustic/phonetic similarity are increasingly well-documented, fundamental issues about exactly how to characterize similarity remain unresolved. These issues are intertwined with questions about how acoustic information is represented in the lexicon, e.g., as sequences of phonetic segments, phonemes or syllables. For example, the most detailed model of competitor similarity is the Neighborhood Activation Model developed by Luce and colleagues (for an overview, see Luce & Pisoni, 1998). In this model, each word is stored as a string of phonetic segments. Lexical activation is a function of overall segment similarity with all other words, weighted by the relative frequency of the target word and the frequency-weighted “neighborhood” within which the target word resides. Similarity is evaluated using the Neighbor Word Probability rule, which provides graded similarity scaling of words in terms of the relative confusability of segments of one word with the other, such that the neighboring words may not necessarily share segments at all. An alternative “shortcut” metric defines a neighbor any word that can be changed to another by adding, deleting, or changing a single segment (neighbors of *cat* include *cap* [a cohort], *hat* [a rhyme], *cot*, *cast*, *at*, but not *catalyze*). Neighbors are then used to compute the neighborhood similarity function that maps onto the processing notion of activation. Strong support for the concept of lexical neighborhoods comes from evidence that words in sparse neighborhoods are recognized more quickly than words in dense neighborhoods.

Despite the strong evidence that positional overlap influences the degree to which a lexical competitor becomes activated, neither of the similarity metrics used by NAM is explicitly sensitive to order information. For example, according to the shortcut rule, a consonant–vowel–consonant (CVC) target word would be equally confusable with other CVC neighbors, regardless of whether the segment difference resided in the initial consonant, the vowel, or the final consonant. (Note that the more fine-grained confusion-based measure of neighborhoods also collapses confusability estimates over syllable position.)¹

Another important question is whether vowels and consonants contribute differently to similarity. There are several reasons why segment type may affect competitor similarity. Vowels are louder, longer, and more robust to noise masking (Dorman, Kewley-Port, Brady, & Turvey, 1977; Horii, House, & Hughes, 1970) than consonants. However, they also vary greatly across talker and phonetic context, and are more subject to diachronic changes than consonants (Hillenbrand, Getty, Clark, & Wheeler, 1995; Peterson & Barney, 1952). In contrast, consonants are distributed more tightly along such dimensions as voice onset time (Lisker & Abramson, 1964; Macmillan, Goldberg, & Braidia, 1988). Consonants also tend to outnumber vowels cross-linguistically, and constitute the majority of word onsets. Note that these two qualifications are simply informational ones: more alternatives means more information can be conveyed, and elements with temporal primacy can do more informational work than, all other things being equal, could as easily be done by elements later in the word. Thus, it is not clear if the greater robustness of vowels outweighs their greater variability compared to consonants.

Some evidence suggests that consonants are more central to lexical representations than vowels. For instance, Van Ooijen and colleagues (Cutler, Sebastián-Gallés, Soler-Vilageliu, & Van Ooijen, 2000; Marks, Moates, Bond, & Stockmal, 2002; Van Ooijen, 1996) have demonstrated that consonant information may be more integral to word identity than vowel information. In their “word reconstruction” task, a nonword is presented which can be turned into a real word by alteration of a single phoneme (e.g., “teeble” → *table* or *feeble*). They found that speakers of English (Van Ooijen, 1996), as well as speakers of Spanish and Dutch

¹ It is of interest that the confusion data used in calculating the NWP come from syllables embedded in white noise (Luce, 1986, as cited in Luce & Pisoni, 1998). Knowing that consonants are more strongly masked than vowels, these noise data (and hence the NWP) may underrepresent the discrimination advantage conferred by consonants, classing words with noninitial overlap as more phonologically similar than they actually are, at least in non-noise contexts.

(Cutler et al.; Marks et al., 2002), are faster and more likely to change a vowel to make a real word, suggesting that vowel information is more “mutable” than consonant information, perhaps because vowels contribute less to lexical identity and lexical organization than consonants. Stated another way, words are more lexically distant when they differ by consonants (*feeble*, *teeble*) than when they differ by vowels (*teeble*, *table*).

However, these experiments were based upon listeners’ knowledge of well-learned words. If the final lexical storage state (that of over-learned words) differs from that of newly learned words, it is not certain that consonants are always so integral to lexical identity. Perhaps earlier in learning, the more acoustically salient parts of the word are more reliable for disambiguating lexical entries.

There is only a sparse literature on how representations change as adults learn novel words (Dumay, Gaskell, & Feng, 2004; Gaskell & Dumay, 2003). In an ingenious study, Gaskell and Dumay presented listeners with nonsense words, such as *cathedruke*, in a phoneme detection task. Afterward, they looked for increments or decrements in activation for real-word neighbors (*cathedral*) by presenting these real words in a lexical decision task. More rapid lexical decisions would indicate phonological priming, while slower (delayed) lexical decisions were taken as evidence of lexical competition (interference) from the new nonsense word. They found that activation changed over the course of several days, with facilitated recognition of *cathedral* in a lexical decision task early on, but inhibition of *cathedral* in the same task on later days. This inhibition implies the establishment of a (semantics-free) lexical entry for *cathedruke*. Later work by Dumay et al. (2004) further clarified these results, finding that lexical competition appeared after 24 h, without further exposure to the nonsense words. The competition appeared earlier when the encoding task was phonological than when it was semantic, suggesting that lack of semantic encoding does not hinder entry into the lexicon. Dumay et al. suggest that the passage of time itself, not further exposure or strengthening of lexical traces, may serve to consolidate adults’ phonological representations of words.

The largest literature on the representation of partially learned words comes from research on word discrimination and word learning in infants and young children. One hypothesis in this developmental literature (Charles-Luce & Luce, 1990, 1995) is that early lexical organization differs qualitatively from mature lexical organization in that, during the earliest phase of learning, words are processed holistically without regard to phonetic-feature conjunctions or linear order, and only gradually does the learner come to favor word-initial information. Thus, an overall acoustic

match between words is likely to be as confusing as (or more confusing than) word-initial overlap. Children, for instance, with smaller-than-adult lexicons, also appear to have more sparsely populated lexical neighborhoods, a good storage method if one’s lexical entries are phonologically underspecified (Charles-Luce & Luce, 1990, 1995).

An alternative hypothesis is that early lexical representations and mature lexical representations are qualitatively identical (Coady & Aslin, 2003; Dollaghan, 1994). Coady and Aslin found that, relative to their smaller vocabularies, children’s neighborhoods were actually denser, not sparser, than would be expected by chance; that is, phonological similarity may even be beneficial in lexical acquisition, rather than engendering confusions among indistinguishable lexical entries. Relatedly, Swingley, Pinto, and Fernald (1999) found that 24-month-olds were “garden-pathed” by words that matched initially and diverged later, but not by rhymes. Swingley and Aslin (2000, 2002) found that 14- to 20-month-old infants were sensitive to single-phoneme changes in familiar target words. Thus, phoneme-level differences are noticed even by early word-learners. Prior to word learning, Jusczyk, Goodman, and Baumann (1999) showed that 9-month-old infants listened preferentially to lists of words sharing features or phonemes at onset, but did not preferentially listen to rhymes. These results indicate, in contrast to a holistic, non-directional model, that there is a word-onset preference quite early in lexical development. All of this evidence suggests that children have well-specified and onset-biased lexical entries rather than holistic ones, even from a relatively early point in acquisition. The foregoing literature on childhood lexicons, of course, may not bear directly on the organization of adult lexicons and the manner in which they grow. The key problem in studying the growth of the adult lexicon centers on gaining control over the frequency and neighborhood structure of newly acquired words.

The paradigm that we will use to investigate differences among consonants and vowels as a function of lexical learning builds upon recent work that uses artificial lexicons to gain precise control over the information structure of the lexicon. Magnuson et al. (2003) examined word frequency, cohort, and rhyme effects by having participants learn 16 novel CVCV labels for 16 black-and-white shapes. They then monitored visual fixations to the shape referents of these words to assess transient confusions of cohort and rhyme competitors. Magnuson et al. found suggestive evidence that as the lexicon becomes better learned, rhyme competition becomes weaker relative to cohort competition. Unlike Charles-Luce and Luce (1990, 1995), they attribute this reduction in rhyme competition not to a change in the nature of lexical storage, but to a progression from weak

to strong lexical traces as more input accrues, with onsets taking on greater predictive value as exposure frequency, and underlying representational fidelity, increases. These results suggest, in line with Coady and Aslin (2003) and Dollaghan (1994), that the process by which the lexicon grows may be very similar in both children and adults.

Importantly for the studies presented here, Magnuson et al. (2003) established two basic results that validate the use of the artificial lexicon paradigm. First, the temporal dynamics of lexical competition in the artificial lexicon were strikingly similar to those for real words, as indicated by frequency effects, neighborhood effects, and the time course of cohort and rhyme competition. Second, the processing of words in the artificial lexicon was largely encapsulated from the native-language lexicon, as assessed by manipulations of neighborhood density between the artificial lexicon and the native-language lexicon (see Magnuson et al., Experiment 3). This latter finding suggests that, at least in the early phase of lexical learning, the artificial lexicon paradigm provides a well-controlled test-bed for examining neighborhood structure without contamination by the native-language lexicon. Not only can we control word frequency across the entire lexicon, but we can also manipulate the phonetic content and the varieties of phonological similarity within the lexicon, venturing outside the normal patterns of most natural languages.

The goal of the present series of experiments was to examine differences in the importance of segment types in lexical representations. To achieve this goal we modified Magnuson et al.'s (2003) artificial lexicon paradigm in two ways. First, we increased the size of the lexicon. Second, and importantly, we extended the use of this paradigm to error patterns, instead of eye fixation proportions, as the dependent measure. While Magnuson et al. trained learners to ceiling performance, we chose to focus on errors as a more sensitive measure of the progression of lexical learning. Using this paradigm, we address two major theoretical points. First, we ask whether representations of newly learned words are more holistic than the phonologically well-specified, onset-weighted representations of overlearned words. If newly learned words show more onset-based confusions than overall-form confusions, hypotheses that suggest qualitative changes over time in lexical entries are cast into doubt. Experiments 1 and 2 evaluate the holistic hypothesis by examining the importance of onset-based similarity versus global similarity, comparing cohort and rhyme confusions and their change with degree of learning. We found that confusions for recently learned words show cohort and rhyme effects similar to those obtained with real words and with a well-learned lexicon. Moreover, onset-based similarity is important

early in learning and does not change with increased learning.

Experiments 3 through 6 then examine the importance of consonants and vowels for lexical storage. Experiment 3 establishes that consonant similarity is a more reliable predictor of confusions than vowel similarity. The remaining experiments evaluate alternative hypotheses for why consonants might contribute more to similarity than vowels. Experiment 4 examines the hypothesis that consonants are more important because they provide more information by using a lexicon with more vowels than consonants, thus rendering the vowels more informative. Experiments 5 and 6 examine whether consonants are more significant than vowels in lexicons in which vowels occur more often in word- and syllable-initial position compared to consonants.

Experiment 1

The goal of Experiment 1 was to determine whether activation of early lexical entries, as reflected in the pattern of errors in matching an auditory word-form to a novel visual shape, would mirror the cohort competitor effects observed in mature lexical activation. Cohort competitors are words that match phonemically at onset (e.g., *candy* and *candle*, as well as *can*, *candor*, *cabbie*, etc.). Such lexical items that match at onset typically produce very strong effects, such as greater response latencies to the target, stronger priming, and larger proportions of eye movements to the competitor than to unrelated items. Magnuson et al. (2003) showed that cohort effects are present in newly learned lexical items as assessed with a lexical activation task and an eye-tracking measure. Their results were interpreted as an indication of lexical similarity between target words and their cohorts, suggesting that cohorts might also be confusable even for very new lexical entries.

Method

Participants

Twelve University of Rochester undergraduates were paid \$30 for three 30 to 40-min sessions on consecutive days. None reported any history of hearing difficulties.

Materials

Participants learned 40 novel words for 40 novel black-and-white objects. Words were CVCV in structure (see Appendix A). All of the pictures had labels, and all of the labels had pictures (that is, no pictures or labels served as meaningless fillers). In one (Cohort) set of 20 items, 10 pairs of words were paired by cohort similarity, defined by an identical initial CVC (*bamo*, *bami*); in a



Fig. 1. On a learning trial, two objects were presented in random locations. After 800 ms, a box formed around one object and its name was spoken.

second set of 20 items, 10 pairs of words were phonologically dissimilar, being paired randomly. The inclusion of pseudo-matched pairs served as an estimate of confusions without phonological similarity.² All words were constructed from a set of 10 consonants ([b, d, g, p, t, k, m, n, s, z]) and 5 vowels ([a, i, e, o, u]). Across all 40 words (balanced across cohort and unrelated sets), each segment occurred in first and second syllable position an equal number of times, with the exception that /o/ occurred nine times in syllable 2, while /u/ occurred only 7 times. (As this small deviation from uniformity was in syllable-final position and could not conceivably modulate cohort size, we do not believe it affected the outcome.) Stimuli were generated with the Macintalk speech synthesis system, using the SpeechSaver application (Singer, D'Oliveiro, 2001) to convert sound samples directly to audio files. The voice used was Victoria (high quality), with rate 150, pitch 54, volume 100, and modulation 7.³

Learning

Participants completed the same learning and test phases on each of three consecutive days. On each trial, two unfamiliar pictures (Fig. 1) were presented simultaneously, and after 800 ms the object of interest was enclosed by a black outline square. After 200 ms, the

name of this object was spoken by the synthetic voice. When the participant clicked the word *next* positioned in the center of the screen, this cycle of object presentation, highlighting, and naming was repeated. Each of the 40 pictures was named on 12 trials in the learning phase of each experimental session, for a total of 480 trials per session. As well, each of those 40 pictures appeared 12 times as the unnamed picture on 12 trials per session. Though the objects were designed to be relatively unfamiliar in appearance, we used four different object/word pairings across participants to minimize spurious coincidences of name and object appearance (e.g., the word *muta* being paired with a cow-like object).

Test

The test phase on each day immediately followed the learning phase and consisted of 80 trials (Fig. 2). Four objects appeared on the screen, and after 100 ms the participant was instructed to (e.g.) “Click on the *muta*.” After a single response on each trial, the next trial, containing four objects, was immediately presented. The positions of the shapes on each trial were counterbalanced so that each shape was equally likely to appear in any of the four screen positions. No feedback was provided during testing.

All 40 pictures served as the target picture on two test trials: once with the competitor picture present (i.e., the *bamo* was the target and the *bami* was also present, as shown in Fig. 2), and once with the competitor picture absent (the *bamo* was the target but the *bami* was not among the objects pictured on the screen). On the competitor-picture-present trial, the

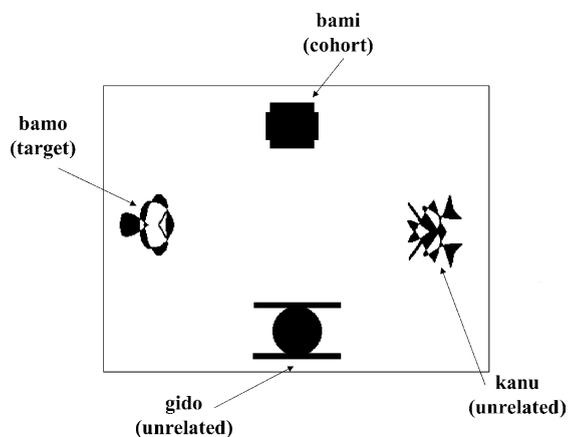


Fig. 2. On a test trial, four objects were presented. On the trials of interest (competitor trials), one object (e.g., *bamo*) was the target object, another (*bami*) a phonologically related competitor (or an unrelated pseudocompetitor), and the remaining two were phonologically unrelated distractors (*gido*, *kanu*).

² In a manipulation that has no bearing on the results of the current experiment, half of each set of pairs contained objects that were always paired with one another in the learning phase, while the other half were randomly paired with a number of items in the learning phase (see Fig. 1). Consistent or random pairing was evenly distributed between cohort and non-cohort pairs. This manipulation was designed to create differences in contextual similarity, which we hypothesized might have effects on confusions in addition to acoustic-phonetic similarity effects. It did not, and we therefore do not further analyze this factor. This method of presentation was maintained in the following experiments only for methodological comparability.

³ Rate is indicated in words per minute, pitch specifies mean F_0 in semitone units, and modulation specifies the range of allowable pitch deviations from the mean (monotone modulation = 0). Volume is expressed as a percentage of the maximum system volume, ranging from 0 to 100.

target picture and the cohort competitor (or pseudo-competitor) picture were accompanied by two pictures with phonologically unrelated names (e.g., **bamo** [target], **bami** [cohort], gido [unrelated], kanu [unrelated]). On the competitor-picture-absent trial, the target was accompanied by *three* pictures with phonologically unrelated names (e.g., bamo [target], kobi [unrelated], gido [unrelated], kanu [unrelated]). This latter trial type was included so that participants would not become aware of the cohort similarity manipulation. Thus, there were 40 competitor-picture-present trials, and these are the ones we examine. Twenty of these trials tested the actual cohort words, the other twenty tested the pseudomatched words. On 20 of these trials the picture of the cohort word was present and selection of that picture when the target word was spoken constituted a cohort confusion. On the other 20 trials, with the other 20 (pseudomatched) targets, the picture of the pseudomatched word was present (e.g., beika [target], gido [pseudocompetitor], tuzi [unrelated], gopu [unrelated]), providing chance confusion data. If cohort confusions occur, there should be relatively more erroneous selections of the cohort picture than of the unrelated pictures on those trials, while the pseudomatched picture should not be selected more often than unrelated pictures.

Analysis

Competitor errors were summed for each participant for each of the three days. Similarly, distractor errors were summed for each participant and then divided by 2 because distractors occurred twice as often as competitors. By comparing rates of competitor errors to rates of distractor errors on each trial type, we could determine whether competitor choices occurred more often than chance, indicating that targets are more lexically similar to cohort competitors than to phonologically unrelated items.

Results and discussion

All of the analyses in this and the following five experiments were restricted to competitor-present trials, and were performed with both participants and items as random factors. We calculated the percentages of errors to the phonologically related competitor (or pseudo-competitor) and the average percentage of distractor errors (recall that *two* unrelated distractor pictures were present on each trial). As shown in Fig. 3, the overall pattern of results showed that error rates to competitors greatly exceeded errors to distractors in the cohort condition. However, error rates for unrelated distractors and for competitors in the pseudomatched condition were comparable to one another. Distractor error rates were similar in both conditions. While the magnitudes of these effects changed across the three days of the

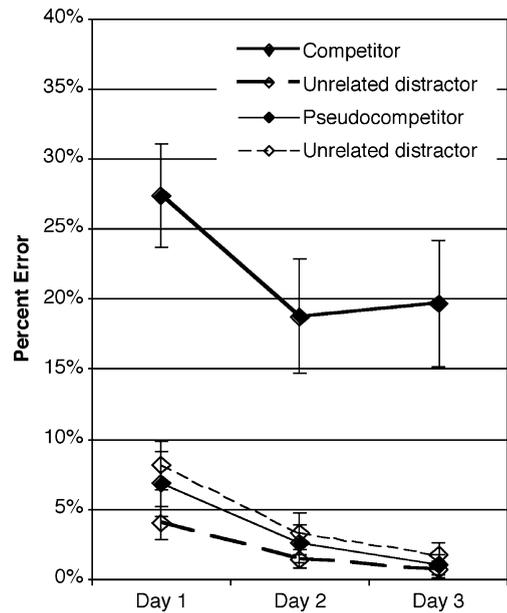


Fig. 3. Experiment 1 results. Thick lines are cohort-trial errors, while thin lines are pseudo-matched trial errors. Solid lines indicate errors to the competitor (or pseudocompetitor), with dashed lines indicating the corresponding unrelated distractor errors divided by 2. Error bars are standard errors.

experiment, the nature of the effect remained similar on all days.⁴

A repeated-measures analysis of variance (ANOVA) with Match Type (cohort or pseudomatched), Error Type (competitor or distractor), and Day (1, 2, 3) as factors confirmed the foregoing summary. There were effects of Day ($F_{1(2,22)} = 10.7$, $p = .0006$; $F_{2(2,76)} = 23.04$, $p < .0001$), indicating decreasing errors over Days; Match Type ($F_{1(1,11)} = 18.72$, $p = .001$; $F_{2(1,38)} = 24.06$, $p < .0001$), indicating overall more errors on the cohort target trials than on pseudomatched target trials; and Error Type ($F_{1(1,11)} = 26.94$, $p = .0003$; $F_{2(1,38)} = 36.41$, $p < .0001$), indicating overall more competitor than distractor errors. There was also a Match Type \times Error Type interaction ($F_{1(1,11)} = 31.48$, $p = .0002$; $F_{2(1,38)} = 44.47$, $p < .0001$), stemming from the fact that competitor and distractor errors only differed in the cohort condition

⁴ In this experiment and the one following, there were no striking effects of or interactions with Day, nor were there such effects in any of the following experiments. There is some interest in the literature as to whether rhyme effects begin stronger and diminish over time as a word becomes more well-learned, while cohort effects remain or even strengthen over learning. This is similar to the developmental debate as to whether words are holistically represented initially, or if they are fully specified from the outset. We find no such effects of stronger rhymes or holistic storage here; rather, rhyme (holistic) confusions are weaker from the outset.

($t_1(11) = 5.59$, $p = .0002$; $t_2(19) = 6.81$, $p < .0001$), and not in the pseudomatched condition ($t_1(11) = -.93$, $p = .37$; $t_2(19) = -.88$, $p = .39$). There were no interactions of any of these factors with Day: Day \times Match Type ($F_1 < 1$; $F_2 < 1$), Day \times Error Type ($F_1(2,22) = 1.09$, $p = .35$; $F_2 < 1$), Day \times Match Type \times Error Type ($F_1(2,22) = 1.57$, $p = .23$; $F_2 < 1$).

These results indicate that partially learned pairs of words sharing acoustic onsets (i.e., first syllable + second syllable onset) were confused with one another more often than with acoustically dissimilar distractor words. This effect was steady over days, while absolute error rates declined. This pattern of results is similar to the lexical activation effects (cohort confusions) seen in experiments on spoken word recognition, using both overlearned words from the natural lexicon (e.g. Marslen-Wilson, 1987) and novel words (Magnuson et al., 2003). Interestingly, our results show that word-initial information is prominent even in partially learned lexical entries, as attested by using error rates—a sign of incomplete mastery of this artificial lexicon—as the dependent measure.

Experiment 2

Experiment 2 further explored whether activation of recently learned lexical items shows the onset biases associated with processing words in a mature lexicon. In contrast to cohort effects, transient lexical confusions of rhymes (words that mismatch only at onset) have been observed only rarely in many tasks (e.g., gating and lexical decision), and when they are reliably present in eye movement tasks, they are much weaker than cohort effects. However, some have argued (Charles-Luce & Luce, 1990, 1995) that for new lexical entries, global acoustic similarity may be more important than the match at word onset, suggesting that rhyme confusions among early lexical entries may be as prevalent as cohort confusions. In the present experiment, we tested this hypothesis by utilizing items which, like the cohort items in Experiment 1, differed by a single segment; but in this case it was the initial segment rather than the final segment, yielding rhyme pairs.

Method

Participants

Twelve University of Rochester undergraduates were paid \$30 for three 30 to 40-min sessions on consecutive days. None reported any history of hearing difficulties and none had participated in Experiment 1.

Materials

All 40 novel words used were CVCV in structure (see Appendix A). In one (rhyme) set of 20 items, 10 pairs of

words shared a final VCV (dutei, mutei), while in a second set of 20 items, 10 pairs of control words did not share final VCVs. Both sets of items contained each segment in both first and second syllable position an equal number of times. All words were constructed from the same set of 10 consonants ([b, d, g, p, t, k, m, n, s, z]) and 5 vowels ([a, i, e, o, u]) used in Experiment 1 and were generated synthetically in an identical manner.

Learning and test

The learning and test phases were in all respects identical to that of Experiment 1, except that instead of cohort pairs, acoustically similar words were rhymes.

Results and discussion

As shown in the right half of Fig. 4, the overall pattern of results showed that errors to rhymes of the target exceeded errors to unrelated distractors, while errors to pseudomatched competitors did not exceed distractor errors. Data from Experiment 1 are presented on the left side of the Figure for comparative purposes; note that the overall cohort effect appears larger than the rhyme effect.

We first analyzed only the data from the current experiment. A repeated-measures analysis of variance (ANOVA) with Match Type (rhyme or pseudomatched) and Error Type (competitor or distractor) as factors confirmed the presence of rhyme confusions: the only significant effect was a Match Type \times Error Type interaction ($F_1(1,11) = 14.6$, $p = .003$; $F_2(1,38) = 7.1$, $p = .01$), showing that competitor and distractor errors only differed in the rhyme condition ($t_1(11) = 2.7$, $p = .02$; $t_2(19) = 2.58$, $p < .02$), and not in the pseudo-

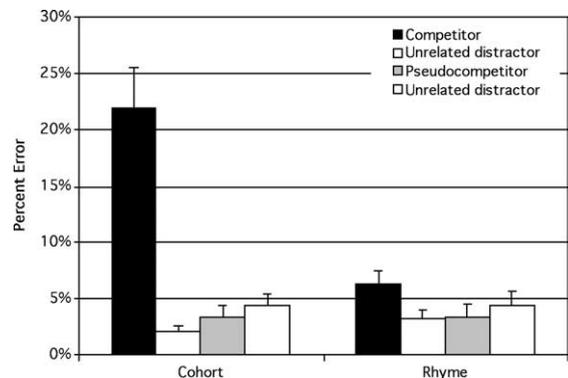


Fig. 4. Experiment 2 (rhyme) results, with Experiment 1 (cohort) results for comparison. Note that Experiment 1 data are now collapsed across Days. Black bars indicate errors to the competitor, gray bars indicate errors to the pseudocompetitor, corresponding white bars indicate corresponding unrelated distractor errors divided by 2. Error bars are standard errors.

matched condition ($t_1(11) = -1.3$, $p = .22$; $t_2(19) = -1.07$, $p = .3$).

Comparing the magnitude of the competitor effect across Experiments 1 and 2 revealed that the cohort effect was greater than the rhyme effect. An Experiment \times Match Type \times Error Type interaction ($F_1(1,22) = 18.74$, $p = .0003$; $F_2(1,76) = 21.93$, $p < .0001$) indicates that the Experiment \times Error Type interaction was significant only for phonological competitors (actual cohorts and rhymes; $F_1(1,22) = 20.44$, $p = .0002$; $F_2(1,38) = 27.45$, $p < .0001$), not the pseudomatched words ($F_1(1,22) = .02$, $p = .89$; $F_2(1,38) = .006$, $p = .94$). That is, there was a greater difference between competitor and distractor errors in the cohort condition than in the rhyme condition. The pseudomatched conditions of both experiments were nearly identical, showing, as expected, no error asymmetries.

These results mirror the asymmetries between rhyme and cohort competition in empirical data on lexical activation (Allopenna et al., 1998; Magnuson et al., 2003; Marslen-Wilson & Zwitserlood, 1989; Marslen-Wilson et al., 1996) and in models of lexical competition such as TRACE (McClelland & Elman, 1986) and Shortlist (Norris, 1994). Our listeners appeared to weight word-initial information more heavily in lexical access, even for incompletely learned lexical entries.⁵ Now that we have established that early lexical entries, much like mature ones, are phonologically well-specified and no more holistically represented than are overlearned entries, we focus on the relative roles of consonants and vowels in lexical representation. In addition, because the pattern of results did not change across days, we collapse our analyses across that variable in order to obtain maximum statistical power for the effects of interest.

Experiment 3

As summarized in the introduction, the literature has provided mixed results on the effects of word similarity defined by consonants and by vowels. Each segment type has potential advantages for carrying useful information—vowels are louder and longer, but are strongly affected by coarticulatory, talker, and random variation (Peterson & Barney, 1952). Van Ooijen and colleagues (Cutler et al., 2000; Marks et al., 2002; Van Ooijen, 1996), by contrast, have shown that consonants may contribute more information than vowels to lexical identity and lexical organization of overlearned words. However, all of Van Ooijen's word similarities were

unidirectional (nonword \rightarrow word), and many contained more consonants than vowels. Here we explore novel word \leftrightarrow novel word confusions, with each word containing an equal number of consonants and vowels.

Method

Participants

Twenty-four University of Rochester undergraduates were randomly assigned to one of two conditions (consonant or vowel similarity) and were paid \$30 per participant for three 30 to 40-min sessions on consecutive days. None reported any history of hearing difficulties, and none participated in Experiments 1 and 2.

Materials

All 40 novel words were CVCV in structure (see Appendix A). In the consonant similarity condition, 10 pairs of words shared two consonants (dozei, duzo), while 10 other pairs of words shared a single initial consonant. In the complementary vowel similarity condition, 10 pairs of words shared two vowels (nasi, tagi), while 10 pairs of words shared the first vowel. All words were constructed from the same set of 10 consonants ([b, d, g, p, t, k, m, n, s, z]) and 5 vowels ([a, i, e, o, u]) used in the previous two experiments. In both conditions, each set of two-segment- or one-segment-matched items contained 5 consonants occurring equally often in the first consonant position, and the other 5 in the second position, with the consonants reversed but also occurring equally often in the other item set. Vowels occurred equally often in each possible position in both sets.⁶ Stimuli were generated synthetically as in Experiments 1 and 2.

Learning and test

The learning and test phases were in all respects identical to that of the preceding experiments, except that two-segment (or one-segment) consonant-(or vowel-) matched items were of primary interest as the potentially confusable words.

⁵ It should be kept in mind that our listeners were all adults, and that first-language acquisition by infants may be a different scenario, a point that will be addressed in the general discussion.

⁶ Because there were a limited number of possible combinations of consonant pairs and vowel pairs, we were forced to use a given consonant pair four times in the experimental words in this and the following experiments: that is, there were four words in the set that had the same pair of consonants or vowels (e.g., pona, punei, pino, panu). However, these quads of items were split into two pairs: each test trial for a given target had only a single competitor. Thus only *pona* and *punei* occurred in the same test trials as target and competitor (or vice versa), and only *pino* and *panu* occurred in the same test trials as target and competitor (or vice versa). This does mean that we cannot directly compare confusion rates between the first two experiments and the final four experiments.

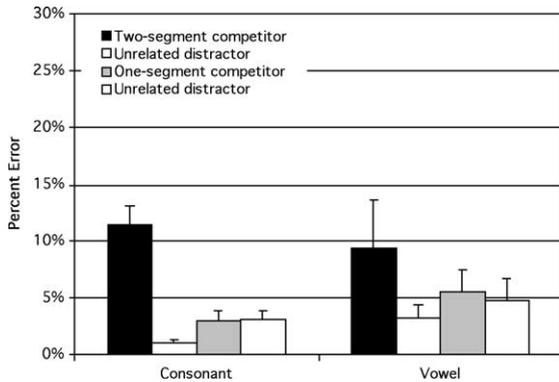


Fig. 5. Experiment 3 results in consonant and vowel conditions. Black bars indicate two-segment-competitor errors, gray bars indicate one-segment-competitor errors, corresponding white bars indicate corresponding distractor errors divided by 2. Error bars are standard errors.

Results and discussion

As shown in Fig. 5, the overall pattern of results revealed more competitor errors than distractor errors for two-segment-matched word pairs, but not for one-segment-matched word pairs. However, this competitor–distractor asymmetry for the two-segment-matched words was only reliable for the consonant condition. In the vowel condition, there appears to be a somewhat smaller (albeit nonsignificant) trend for two-segment-matched words.

A mixed repeated-measures ANOVA with Segment Type (consonant, vowel), Overlap (two-segment, one-segment), and Error Type (competitor, distractor) as factors confirmed the foregoing summary.⁷ The effects of Overlap ($F_1(1,22) = 10.46$, $p = .004$, $F_2(1,33) = 5.73$, $p = .02$) and Error Type ($F_1(1,22) = 8.42$, $p = .008$, $F_2(1,33) = 12.45$, $p = .001$) were significant, as well as the interaction of Overlap and Error Type ($F_1(1,22) = 18.23$, $p = .0003$, $F_2(1,33) = 10.64$, $p = .003$). The Overlap effect indicated more errors on two-segment matched trials than on one-segment matched trials, and the Error Type effect suggested more competitor than distractor errors overall. The interaction of these two factors indicated that the Error Type effect (more competitor

than distractor errors) was stronger for two-segment-matched trials. The three-way interaction of Segment Type, Overlap, and Error Type did not reach significance ($F_1(1,22) = 3.13$, $p = .09$, $F_2(1,33) = 2.36$, $p = .13$).

Because we expected that effects between the two segment types might differ, we performed analyses for each Segment Type condition individually. For the consonant condition, Error Type ($F_1(1,11) = 29.57$, $p = .0002$, $F_2(1,14) = 6.26$, $p = .03$) and the Overlap \times Error Type interaction ($F_1(1,11) = 27.58$, $p = .0003$, $F_2(1,14) = 7.04$, $p = .02$) were both significant, the latter indicating that only for two-segment-matched words did competitor errors exceed distractor errors ($t_1(11) = 6.73$, $p < .0001$, $t_2(9) = 3.42$, $p = .008$).

For the vowel condition, however, the same effects did not reach significance: Overlap, $F_1(1,11) = .84$, $p = .38$, $F_2(1,19) = 1.21$, $p = .29$; Error Type, $F_1(1,11) = 1.23$, $p = .29$, $F_2(1,19) = 5.47$, $p = .03$; Overlap \times Error Type, $F_1(1,11) = 2.34$, $p = .15$, $F_2(1,19) = 2.53$, $p = .13$. These nonsignificant effects suggest that the weight of the Error Type and Overlap \times Error Type effects were carried by the consonant condition, perhaps due to greater variability (lower reliability) in the vowel condition.

The overall pattern of results in Experiment 3 revealed a possible asymmetry in lexical confusions: matching consonant frames are more reliably confusable than matching vowel frames. This is not simply the effect of the two-consonant items being one-consonant cohorts, as words that shared *only* a single consonant were not confused with one another. Words that had to be disambiguated by their vowels (two-consonant-matched words) were more reliably confused with one another than words that had to be disambiguated by their consonants (two-vowel-matched words). It should be noted here that the statistically significant two-consonant confusions involved items that differ by two segments, suggesting that the shortcut rule for estimating neighborhood size in NAM (Luce & Pisoni, 1998), which has been used extensively in the literature on lexical activation, may provide an incomplete metric for neighborhood similarity.

The present results suggest a tendency for consonants to be somewhat more integral to word identity than vowels. However, whether this asymmetry is actually present is unclear, and at least two alternative explanations should be considered before accepting this conclusion. First, it may be the case that consonants are simply less altered by coarticulation than vowels, making consonant information slightly more reliable than vowel information in word identification. Second, and perhaps more obviously, in our set of words, consonants outnumbered vowels by a 2:1 ratio. Thus vowels were less indicative of word identity, possibly obscuring effects of similarity and accounting for the nonsignificant

⁷ Due to a design oversight in Experiments 3 and 4, a large proportion of trials contained distractors which shared one (or more) initial C or V with the target item. When the data from these trials were removed, we found that the results were qualitatively similar to those in the full analysis, while lower in power (analyses reported are from the reduced analysis). Importantly, in both experiments, consonant confusions were still more statistically robust than vowel confusions.

result. In the next experiment we reverse this ratio to determine whether relative segment number within an artificial lexicon can modulate word confusions.

Experiment 4

The segment inventory in all of the preceding artificial lexicon experiments consisted of more consonants than vowels, a distributional property of English and many other natural languages. The greater number of consonants rendered them more informative as to lexical identity than vowels. While our previous stimuli mirrored segment type distributions in most human languages, it is possible to create an artificial lexicon that does not have these properties. In the present experiment we use artificial lexical items constructed from a different set of segments having a greater number of vowels than consonants. If informativeness within the artificial lexicon is important in the lexical weighting of segment types, then vowel-based confusion effects should emerge as the relative size of the vowel inventory exceeds that of consonants.

Method

Participants

Twelve University of Rochester undergraduates were randomly assigned to the consonant similarity condition and 12 to the vowel similarity condition, and were paid \$30 per participant for three 30 to 40-min sessions on consecutive days. None reported any history of hearing difficulties, and none was tested in Experiments 1–3.

Materials

All 40 novel words were CVCV[C] in structure (see Appendix A). An additional non-informative consonant⁸ (hence the notation [C]), d₃, was placed at the end of each word to allow the use of lax vowels. In the consonant similarity condition, 10 pairs of words shared the first two consonants (bosaid₃, basəd₃), while 10 other pairs of words shared a single initial consonant. In the complementary vowel similarity condition, 10 pairs of words shared two vowels (tɪbaɪd₃, pɪgaɪd₃), while 10 pairs of words shared the first vowel.

In contrast to the preceding experiments, all words in the present experiment were constructed from a set of 5 consonants ([b, g, p, t, s]) and 10 vowels ([a, i, eɪ, o, u, aɪ, əɪ, aʊ, ɪ, ε]), such that vowels now outnumbered consonants by a ratio of 2 to 1. In both conditions, each set of

two-segment- or one-segment-matched items contained 5 vowels that occurred equally often in the first vowel position, and the other 5 vowels in the other position, with the vowels reversed but also occurring equally in the other item set. Consonants occurred equally often in first and second position in both sets. Stimuli were generated synthetically as in the previous three experiments.

Learning and test

The learning and test phases were in all respects identical to that of the preceding experiments.

Results and discussion

As shown in Fig. 6, the overall pattern of results revealed that in the consonant condition there were competitor effects for both two-segment-matched and one-segment-matched words, suggesting robust consonant-based confusions. However, unlike Experiment 3, there was not a numerical tendency toward vowel-based confusions. It seems that simply altering the relative proportions of segment types in the artificial lexicon is insufficient to improve the capacity of vowels to disambiguate words.

A mixed repeated-measures ANOVA with Segment Type (consonant, vowel), Overlap (two-segment, one-segment), and Error Type (competitor, distractor) as factors confirmed the foregoing summary. The effect of Error Type was significant ($F_1(1,22) = 7.82, p = .01, F_2(1,32) = 9.13, p = .005$), indicating more competitor than distractor errors overall. The Segment Type \times Error Type interaction ($F_1(1,22) = 7.85, p = .01, F_2(1,32) = 8.83, p = .006$) also reached significance, due to a greater competitor–distractor difference for the consonant condition. The overall Segment Type \times Overlap \times Error Type interaction did not approach significance ($F_1(1,22) = .94, p = .34, F_2(1,32) = .92, p = .34$).

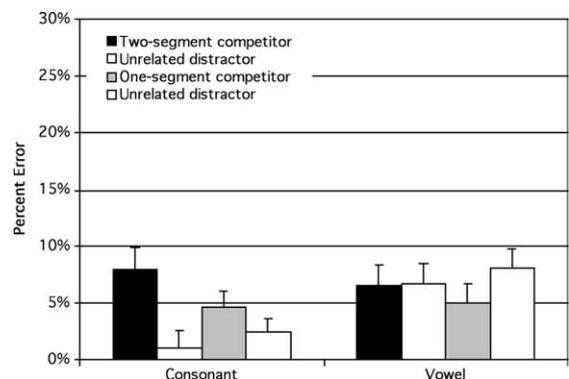


Fig. 6. Experiment 4 results in consonant and vowel conditions. Black bars indicate two-segment-competitor errors, gray bars indicate one-segment-competitor errors, corresponding white bars indicate corresponding distractor errors divided by 2. Error bars are standard errors.

⁸ In Experiment 4 and later in Experiment 6, we choose to utilize a single consonant that is redundant for all words. That is, “non-informative” indicates that identifying this consonant is unhelpful in distinguishing among any pair of words in the artificial lexicon.

We then examined each Segment Type condition individually. For the consonant condition, the effect of Error Type was significant ($F_1(1,11) = 11.96, p = .005, F_2(1,16) = 13.34, p = .002$), but not the Overlap \times Error Type interaction ($F_1(1,11) = .31, p = .59, F_2(1,16) = .28, p = .6$). While the two-segment items ($t_1(11) = 2.33, p = .04, t_2(9) = 2.78, p = .02$) showed a significant competitor effect, the one-segment competitor effect did as well ($t_1(11) = 3.53, p = .005, t_2(7) = 2.6, p = .04$), though slightly smaller in magnitude (note that this was not a significant difference). Thus, competitor errors were greater than distractor errors when the competitor shared at least one consonant with the target. However, for the vowel condition, no effects approached significance: Overlap: $F_1(1,11) = 2.03, p = .18, F_2(1,16) = 1.00, p = .33$; Error Type: $F_1(1,11) = 0, p = 1, F_2(1,16) = .002, p = .97$; Overlap \times Error Type: $F_1(1,22) = .74, p = .41, F_2(1,32) = .86, p = .37$. Thus, same-vowel confusion effects went undetected in this experiment.

These results indicate that both single- and double-consonant-matched items are more reliably confused with one another than are vowel-matched items (recall the significant Segment Type \times Error Type interaction). Nonetheless, vowels do not seem to be more informative than in the preceding experiment, even when they vastly outnumber consonants. This elaborates upon the preceding experiment, implying that vowel information is not weighted more highly when the inventory of vowels is larger and vowels are, by definition, more informative for word identity than consonants. Indeed, it may be the case that having a large number of vowels makes them *less* useful as cues to word identity, perhaps because a larger number of vowels may increase confusability as they begin to overlap in vowel space. It is also possible that, while vowels would otherwise be informative, speakers of American English tend to downweight vowel information in lexical entries because vowels are less stable than consonants across dialects. However, the only way to determine whether vowel information can take on added importance in lexical representation would be to greatly extend the duration of the learning phase for an artificial lexicon that is distributionally “vowel heavy” as in Experiment 4.

Experiment 5

The fact that consonant similarity still led to more detectable confusability than vowel similarity, at least after 60 to 90 min of exposure to an artificial lexicon where vowels outnumbered consonants, suggests that consonants may hold a privileged status in lexical representations. However, consonants in the preceding experiments possessed another advantage which we have yet to confer upon the vowels: word-initial position. Perhaps words with similar consonants are more confusable

partly due to the fact that most words in English (and many other languages) begin with a consonant (or consonant cluster). According to this view, consonants take on added weight in lexical activation because they differentiate words (or narrow the range of alternative words) immediately after the processing of the initial phonetic segment (i.e., they define a cohort), and this consonant-bias spreads to later-occurring consonants in the word. To test this hypothesis, we created artificial lexicons that consist entirely of VCVC words. If word-initial position is the major contributor to the consonant confusion effects we have seen in the previous experiments, then we should find reliable evidence of vowel confusion effects and perhaps a reduction in consonant confusion effects.

Method

Participants

Twelve University of Rochester undergraduates were randomly assigned to a consonant similarity condition and 12 to a vowel similarity condition, and were paid \$30 per participant for three 30 to 40-min sessions on consecutive days. None reported any history of hearing difficulties, and none was tested in Experiments 1–4.

Materials

The 40 novel words were VCVC in structure (see Appendix A). In the consonant similarity condition, 10 pairs of words shared two consonants (opan, upein), while 10 other pairs of words shared the first consonant. In a complementary vowel similarity condition, 10 pairs of words shared two vowels (umeis, upei**b**), while 10 pairs of words shared the first vowel.

As the results of Experiment 4 did not suggest that increased vowel number raised participants’ weightings of vowel information, all words in the present experiment were constructed from the original set of 10 consonants and 5 vowels. In both conditions, each set of two-segment- or one-segment-matched items contained 5 consonants occurring equally often in the first consonant position, and the other 5 in the second consonant position, with the consonants reversed but also occurring equally often in the other item set. Vowels occurred equally often in each of the two positions in both sets. Stimuli were generated synthetically as in the previous experiments.

Learning and test

The learning and test phases were in all respects identical to that of the preceding experiments.

Results and discussion

As shown in Fig. 7, the overall pattern of results revealed that there were both consonant-based and vow-

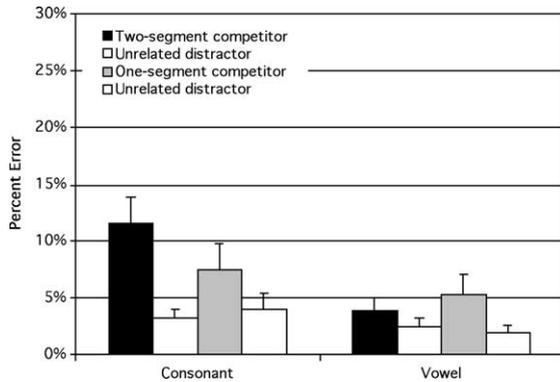


Fig. 7. Experiment 5 results in consonant and vowel conditions. Black bars indicate two-segment-competitor errors, gray bars indicate one-segment-competitor errors, corresponding white bars indicate corresponding distractor errors divided by 2. Error bars are standard errors.

el-based confusion errors, but the consonant based ones were greater. Both two-consonant and one-consonant competitors were selected more often than distractors on the same trials, and these error rates (two-segment and one-segment) were not different from one another. For vowels, interestingly, only the one-vowel confusion effect reached significance. These results suggest that consonant effects were mainly second-syllable-onset (CVC) effects, while the vowel effects were first-syllable-match (VC) effects, with the consonant (second-syllable-onset) effects being greater in magnitude.

A mixed repeated-measures ANOVA with Segment Type (consonant, vowel), Overlap (two-segment, one-segment), and Error Type (competitor, distractor) as factors confirmed the foregoing summary.⁹ The effect of Error Type ($F_1(1,22) = 21.16$, $p = .0001$, $F_2(1,72) = 21.02$, $p < .0001$) and the Segment Type \times Overlap \times Error Type interaction ($F_1(1,22) = 11.84$, $p = .002$, $F_2(1,72) = 3.99$, $p < .05$) reached significance. As usual, the Error Type effect indicated more competitor errors than distractor errors overall. To examine the interaction, we then analyzed each Segment Type condition individually.

For the consonant condition, the effect of Error Type was significant ($F_1(1,11) = 14.59$, $p = .003$, $F_2(1,35) = 13.04$, $p = .0009$); both two-segment ($t_1(11) = 4.04$, $p = .002$, $t_2(17) = 2.77$, $p = .01$) and one-segment competitor effects ($t_1(11) = 2.47$, $p = .03$, $t_2(18) = 2.4$, $p = .03$) were significant. The Overlap \times Error Type interaction only reached significance by participants ($F_1(1,11) = 7.47$, $p = .02$, $F_2(1,35) = 1.94$, $p = .17$). That is, two-

consonant-matched items and one-consonant-matched items were confused with one another, with the two-consonant effect being numerically greater.

For the vowel condition, Error Type was significant ($F_1(1,11) = 6.61$, $p = .03$, $F_2(1,37) = 9.06$, $p = .005$), but the Overlap \times Error Type interaction was not ($F_1(1,11) = 4.75$, $p = .052$, $F_2(1,37) = 3.01$, $p = .09$). However, only one-segment errors exceeded chance ($t_1(11) = 2.58$, $p = .03$, $t_2(19) = 3.43$, $p = .003$). In contrast, the competitor–distractor difference for two-segment errors missed significance ($t_1(11) = 2.15$, $p = .055$, $t_2(18) = .88$, $p = .39$), though in the correct direction. An important fact here is that double-vowel confusions do not exceed single-vowel confusions, indicating that initial-vowel similarity was sufficient to drive confusions (i.e., the second vowel apparently did not increase similarity).

The results do not unequivocally support the hypothesis that word-initial position is necessary for a segment type to elicit confusions. On the one hand, consonant confusion effects are still present. On the other hand, a vowel effect is present, if only weakly so. In addition, the fact that initial cohort confusions with vowels, which are robust when words begin with consonants, did not exceed confusions of words mismatched at onset (two-consonant-matched), provides further support for the importance of consonants.

Despite the fact that the present results suggest a reduced weight for vowel information, even when vowels begin every word in the artificial lexicon, two other factors may be important. First, as in Experiment 4, it is possible that with additional exposure, any effects of English (where consonants begin most words) would be reduced or eliminated, thereby revealing a robust up-weighting of vowel information. Second, while vowels in the present experiment were word-initial, the second vowel was always in syllable-medial position, unlike the consonants in Experiments 3 and 4. This is because in English phonology, generally speaking, a VCVC word is syllabified as V.CVC, whereas a CVCV word is syllabified as CV.CV (though see the General Discussion). Thus the consonant advantage might be a syllable-initial advantage rather than a word-initial advantage.

Experiment 6

In this final experiment, we explore whether the locus of the consonant confusion effect is at syllable onset. It is not possible within the English phonological system to create two-syllable words in which the second vowel would be unambiguously perceived as the onset of the second syllable. However, it is possible to create two-syllable words where the first *disambiguating* element in each syllable is a vowel, by giving the words a VC[C]VC

⁹ In Experiments 5 and 6, a small number of test trials (3 in the consonant condition, 1 in the vowel condition) were eliminated from the analyses because two identical distractor items appeared, changing the probability that each would be chosen.

structure (e.g., by inserting a single segment—here, /f/—after the first consonant of a VCVC). This forces each word to syllabify as VC.[C]VC.¹⁰ Thus, we created two-vowel items that are matched at syllable onset using this VC[C]VC structure. If previous consonant effects were due to their syllable-initial position, vowel-initial effects should emerge in the present artificial lexicon.

Method

Participants

Twelve University of Rochester undergraduates were randomly assigned to the consonant similarity condition and 12 to the vowel similarity condition, and were paid \$30 per participant for three 30 to 40-min sessions on consecutive days. None reported any history of hearing difficulties, and none was tested in Experiments 1–5.

Materials

The 40 novel words were VC[C]VC in structure (see Appendix A). Words were identical to those in Experiment 5, except that a non-informative consonant /f/ (notated [C]) was inserted into each word after the first consonant, to force the syllabification VC.[C]VC, causing all vowel-matched words to be identical at their syllable onsets. We selected the segment /f/ for a number of reasons. First, it was not contained within the set of consonants from the Experiment 5 stimuli, which we hoped to modify only minimally for maximum comparability. Second, it does not participate in many onset clusters in English, none of which occurred in our stimuli. Moreover, it participates in only one onset cluster as the second element, (/s/, as in /sfir/). It also occurs in relatively few coda clusters (only /mf/ in our set) in English phonology. By using this segment, we changed the syllabification of the first consonant in the Experiment 5 stimuli

such that now the informative consonants were all syllable-final, while the first informative element in each syllable was a vowel.

In the consonant similarity condition, 10 pairs of words shared the first and third consonants (opfan, upfein), while 10 other pairs of words shared the first consonant. In a complementary vowel similarity condition, 10 pairs of words shared two vowels (umfeis, upfeib), while 10 pairs of words shared the first vowel. All stimuli were generated synthetically as in the previous five experiments.

Learning and test

The learning and test phases were in all respects identical to that of the preceding experiments.

Results and discussion

As shown in Fig. 8, the overall pattern of results suggested for the first time in this series of experiments that vowel confusion effects were substantially larger than consonant confusion effects. However, this vowel effect only surfaced in the two-vowel case and not in the one-vowel case, suggesting that the two-vowel-matched competitors were more confusable than the one-vowel-matched competitors.

A mixed repeated-measures ANOVA with Segment Type (consonant, vowel), Overlap (two-segment, one-segment), and Error Type (competitor, distractor) as factors confirmed the foregoing summary. The effect of Error Type was significant ($F_1(1,22) = 8.83, p = .007, F_2(1,72) = 10.57, p = .002$), indicating more competitor errors than distractor errors overall. Also significant were the interactions of Segment Type \times Overlap ($F_1(1, 22) = 5.71, p = .03, F_2(1,72) = 6.53, p = .01$), Overlap \times Error Type ($F_1(1,22) = 9.05, p = .007, F_2(1,72) = 5.29,$

¹⁰ To document how our participants were assigning the first consonant to a given syllable, we conducted a control experiment in which participants listened to the words from Experiments 5 and 6 and pronounced these VCVC or VCfVC stimuli with a break in between syllables. Productions were scored by a naive rater according to the assignment of the first consonant to the syllable containing the preceding vowel or the syllable containing the following vowel. For VCVC stimuli, 90% of the first consonants were assigned to the syllable containing the second vowel (CVC); another 4% were ambisyllabic, and only 3% were assigned solely to the syllable containing the preceding vowel (the remaining 3% were indeterminate). For VCfVC stimuli, 94% grouped the first consonant with the first vowel and treated the fVC as a syllable (the remaining 6% were indeterminate). These results suggest that learners in Experiment 5 placed the syllable boundary between the first vowel and the first consonant, making it an onset-consonant, while the learners in Experiment 6 placed the syllable boundary after the first consonant, making it a coda-consonant.

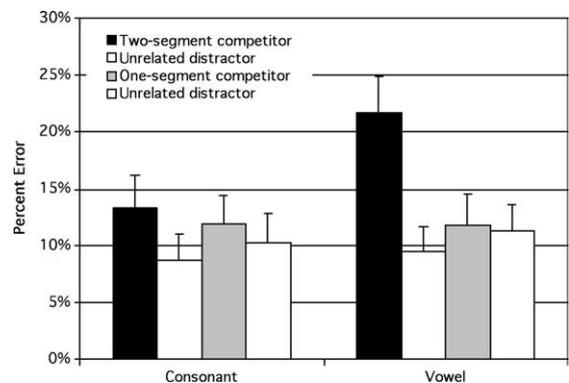


Fig. 8. Experiment 6 results in consonant and vowel conditions. Black bars indicate two-segment-competitor errors, gray bars indicate one-segment-competitor errors, corresponding white bars indicate corresponding distractor errors divided by 2. Error bars are standard errors.

$p = .02$), and Segment Type \times Overlap \times Error Type ($F_1(1,22) = 14.31$, $p = .001$, $F_2(1,72) = 9.71$, $p = .003$). This last interaction was clarified by the individual Segment Type condition analyses.

For the consonant condition, no effects approached significance, indicating the absence of detectable confusion effects. Overlap ($F_1(1,11) = .90$, $p = .36$; $F_2(1,35) = 1.41$, $p = .24$), Error Type ($F_1(1,11) = .56$, $p = .47$; $F_2(1,35) = .20$, $p = .66$), and the Overlap \times Error Type interaction ($F_1(1,11) = .85$, $p = .38$; $F_2(1,35) = .38$, $p = .54$) were all nonsignificant.

However, for the vowel condition, Overlap ($F_1(1,11) = 4.9$, $p < .05$, $F_2(1,37) = 5.59$, $p = .02$), Error Type ($F_1(1,11) = 10.01$, $p = .009$, $F_2(1,37) = 15.79$, $p = .0003$), and the Overlap \times Error Type interaction ($F_1(1,11) = 14.0$, $p = .003$, $F_2(1,37) = 13.28$, $p = .0008$) were all significant. The Overlap effect indicated more errors overall in the two-vowel-matched trials, while the Error Type effect indicated more competitor than distractor errors. The interaction of Overlap \times Error Type was brought about by the fact that only for two-segment-matched words did competitor errors exceed distractor errors ($t_1(11) = 3.67$, $p = .004$, $t_2(18) = 4.58$, $p = .0002$), with no differences for one-segment-matched words ($t_1(11) = .48$, $p = .64$, $t_2(18) = .29$, $p = .78$). Thus, unlike the preceding experiment, a match on two vowels was substantially more confusing than a match on one vowel.

These results provide the first evidence of a robust vowel-confusion effect. It appears that the beginnings of syllables, whether C_C_ or V_fV_, are strongly weighted in newly learned lexical representations, at least when all words have the same syllabic template. Moreover, the consonant effect observed in all of the previous experiments is absent in the present experiment, suggesting that whatever consonant weighting is due to exposure to English can be reduced substantially after only 90 to 120 min of exposure. It should be noted that this apparent down-weighting of consonants may be due to the fact that coda consonants are restricted in their distributions, both in English and cross-linguistically (cf., Ladefoged, 2004), relative to onset consonants. That is, coda consonants are less informative as to word identity because they constitute a smaller set, and they thus discriminate between fewer alternatives than the larger set of onset consonants. Perhaps an even more likely explanation is that coda consonants are less well perceived than onset consonants (Horii et al., 1970; Redford & Diehl, 1999), making them a less reliable source of information for discriminating among lexical alternatives. This is true, for instance, in the case of speech embedded in noise, where consonants become more difficult to perceive (Horii et al., 1970; Johnson, 2000), with confusions retaining (reliable) vowel content while erroneously substituting consonantal content.

Interestingly, the vowel confusion effect is equivalent in magnitude to the comparable consonant confusion effect in Experiment 3. These experiments were structured identically in terms of numbers of neighbors for each word (see Footnote 6), numbers of consonants and vowels used, and the task employed; the only difference was that consonants occupied syllable-initial position in Experiment 3 whereas vowels did in Experiment 6. Accordingly, we compared the consonant condition from Experiment 3 with the vowel condition in Experiment 6 in a repeated-measures ANOVA with Segment Type, Overlap, and Error Type as factors. There were no interactions of Experiment with any other factors ($p \geq .57$). This suggests that differences between competitor and distractor confusion rates were roughly equivalent whether words shared two syllable-initial consonants or two syllable-initial vowels. There was also a strong main effect of Segment Type ($F_1(1,22) = 13.63$, $p = .001$; $F_2(1,51) = 35.11$, $p < .0001$), with more errors overall for the matched-vowel subjects. This overall elevation of error rates may betray an underlying consonant-confusion effect — all words in Experiment 6 shared a single syllable-initial consonant (/f/) — or simply a greater difficulty in learning words that are somewhat atypical in their (roughly) vowel-initial syllables. We address this further in the General Discussion.¹¹

General discussion

In this series of studies we have made several interesting discoveries regarding how segment type and order affect the organization of a partially learned novel lexicon. First, we found that incompletely learned cohort (onset) competitors were more often confused (Experiment 1) than incompletely learned rhyme (offset) competitors (Experiment 2), with cohort confusions far exceeding rhyme confusions, even though both cohorts and rhymes differed from one another by a single (final

¹¹ It has been suggested that the results we interpret as lexical confusions are actually confusions at the perceptual level. Accordingly, we asked 16 participants to judge whether two stimuli were the same or different, using pairs from Experiment 6. Half heard the words from the vowel condition, the other half from the consonant condition. Sometimes, the pair shared vowels (or consonants), and sometimes it did not. For the shared-C and shared-V pairs, percentage correct detections of a difference for items sharing two segments were 100% and 98.3%, respectively. Overall percentages correct for pairs of items that shared or did not share Cs or Vs were 99.9% and 96.9%, respectively. Neither of these results differed significantly by shared-segment type, nor did they differ from 100% correct. These results indicate that participants had little difficulty differentiating items that shared vowel or consonant pairs. Thus, it is unlikely that perceptual-level confusions contributed to our error data.

or initial) segment. This result suggests that the temporal unfolding of a word is important for accessing partially learned lexical entries, not just for the on-line processing of well-learned words. Second, shared consonants in word- and syllable-initial position yielded reliable confusion effects, while vowels in CVCV words did not (Experiments 3 and 4), even when vowels outnumbered consonants, making them more informative as to word identity (Experiment 4). Third, creating VCVC words in which consonants no longer appear in their canonical word-initial position does not substantially elevate vowel confusions (Experiment 5). Finally, words that share vowels are strongly confused (and words that share consonants are not) if all syllables containing vowels are matched at their beginnings (onset + nucleus; Experiment 6).

Our results suggest that partially learned lexical entries may be qualitatively similar to well-learned lexical entries. We found strong cohort effects and weak rhyme effects, reflective of the lexical processing literature, even when errors were used as the dependent measure. Thus, for adults, in a task with a familiar phonetics and phonology, lexical access for novel words and lexical access for well-learned words seem to utilize similar lexical representations. Our findings may also be applicable to vocabulary acquisition during development, insofar as the native phonological system typically solidifies well before vocabulary acquisition reaches asymptotic levels. Thus it will be important to conduct comparable studies with infants and young children. Importantly, our results using a non-time-dependent task mirror those found in time-locked (eye movements) or speeded (lexical decision) tasks.

The overall pattern of results from our experiments suggests that lexical organization is strongly onset-biased. However, our data indicate that the onsets that listeners seem biased to weight strongly in lexical representation are syllable onsets, rather than just word onsets. In practice, this syllable-onset bias typically involves consonant information because that is the prevailing informational structure of English and many other natural languages. However, as we saw in Experiment 6, when syllables are consistently vowel-initial, the syllable onset-bias can switch to match this new information structure to some extent. Taken together, these results hint that lexical representations reflect knowledge of the distributional properties of English, in that cohort effects are the result of both consonant and onset biases given the prevalence of consonant-initial words, while rhyme effects are weaker because competitors are easily distinguished by their consonants.

Why this onset bias exists remains an interesting question for future research. There are several related possibilities. Some explanations are informational in nature, while others are acoustic. It could be, for instance, that elements with temporal primacy (i.e.,

onset consonants) allow for the most rapid information acquisition. Another possibility is that later information (nucleus, coda) is less informative because there are fewer alternatives, leading to an onset bias in word recognition. A more acoustic variant of this hypothesis is that later information (vowels, final consonants) is overall less discriminable acoustically, leading to fewer alternatives *because* a smaller set of alternatives can be discriminated from one another. From an auditory perception standpoint, some argue that it is the case that syllables are the primary objects of speech perception (Mehler, Dommergues, Frauenfelder, & Segui, 1981). More recent work, however, suggests that the average amount of information coming from a single source is a more important predictor of successful word perception (Samuel, 1991): alternating the signal between ears at syllable boundaries versus in mid-syllable did not improve intelligibility in a signal-alternation task. Note that alternating in mid-syllable keeps syllable onsets in the same ear. Consistent with this possibility, originally raised by Samuel, Content, Kearns, and Frauenfelder (2001) argue based on experiments with syllabification word games in French that syllable onsets are easily detected while offsets are not, and thus onsets constitute sensible points at which to initiate lexical access.

These results also have important implications for similarity metrics used in models of lexical neighborhoods. Consider again the two metrics discussed by Luce and Pisoni (1998): the one-segment shortcut rule and the graded Neighbor Word Probability rule. First, while the one-segment rule for defining a word's neighbors may be a useful heuristic for determining neighborhood size, it is inadequate to explain our lexical confusion effects: many confused words (Experiments 3–6) differed by two segments. Too, it does not predict that cohorts differing by one segment (Experiment 1) are more lexically confusable than rhymes that also differ by one segment (Experiment 2). The graded rule may be a more accurate metric of neighborhood structure (but see Bailey and Hahn, 2005; for evidence against it). As Magnuson et al. (2003) note, such a rule would need to take into account order effects: our findings clearly indicate that not all segment positions contribute equally to lexical confusability.

It may also be the case that the graded rule is less generalizable speech that is not embedded in noise: our results suggest that syllable-initial segments may be better indicators of lexical identity, but the graded similarity measure is based on segment confusions from noise-embedded syllables, and this may tend to overemphasize the confusability of consonants (see Footnote 1). Bailey and Hahn (2005) showed that theoretically motivated feature-based models of word similarity predict lexical confusions much better than do empirical measures, including Luce's (1986) confusion data used to calculate Neighbor Word Probability. In addition to this, our results suggest that syllable position should

be accounted for in a model of similarity. Bailey and Hahn's work is useful in that it provides an easily interpreted set of metrics that can be tested against (and possibly refined by) carefully designed stimulus sets that consider segment type and syllable position.

Perhaps the most interesting finding from our series of experiments was that words sharing syllable-initial consonants were more readily confusable than words sharing nuclear vowels, while words sharing syllable-initial vowels were more confusable than words sharing coda consonants. The most important fact here is that the strongest influence on word confusions is a match at the beginning of a syllable. Thus, segment type may not be the critical factor, but rather where the segment tends to occur in a syllable.

On the other hand, one piece of evidence argues for a particularly strong role for consonants, in line with Van Ooijen and colleagues (Cutler et al., 2000; Marks et al., 2002; Van Ooijen, 1996). If one takes into account the fact that there was only one syllable-initial consonant (f) in the entire set of words in the syllable-initial vowel study (Experiment 6), one might conclude that the heightened error rate found in this Experiment was due to massive syllable-onset-consonant confusions, despite the noticeable vowel confusions that were also present. On the other hand, this might be due to the fact that listeners know relatively fewer words that are disambiguated primarily by their vowels (i.e., those words are phonotactically improbable), and these are thus harder to learn overall, with the confusion effect still being equal in magnitude to the consonant-confusion effect. Alternately, vowels might be weighted less strongly for listeners prior to experimental exposure because, in English (as well as most other languages), consonant information is generally word- and syllable-initial, and this more rapidly available information may therefore become more heavily weighted over a lifetime of exposure than vowel information, which becomes available later in the syllable. Thus, the words may simply have been more difficult to learn due to language experience outside the lab.

Other explanations for differences in consonant vs. vowel effects can be ruled out. In Experiment 4, we excluded within-set segment number as a factor, showing that even when vowels outnumber syllable-initial consonants, they do not generate more vowel-based confusions. These results are consistent with the Cutler et al. (2000) finding that word reconstruction effects were no stronger in Castilian Spanish—with 5 vowels and 20 consonants—than in British English, with 17 vowels and 24 consonants (Maddieson, 1984). Another concern was that dialectal variation in English tends to be vowel-based, and thus vowels as cues might be less useful for word identity specifically in English. But again, Cutler et al. showed that effects surfaced in Dutch and Castilian Spanish, both of which have primarily consonantal vari-

ation in dialect. More obviously, we found strong vowel effects when vowels were the first element to identify a syllable. Lastly, it could be objected that our use of synthesized speech might not generalize to natural language. However, several facts argue against this: we obtained results consistent with the lexical processing literature; van Ooijen and colleagues found related effects using natural speech; we obtained confusion effects for both segment types, which would not be the case if, for instance, one segment type were more (randomly or contextually) variable in our synthesized speech than in natural speech.

Finally, the present findings may have implications for lexical development. As mentioned earlier, children's lexicons may well be organized in the same manner as our artificial lexicons in adults; that is, after the phonology of the language has been well established. However, in other respects, the weak vowel-confusion results in Experiments 3 and 4 are not entirely consistent with related developmental work. Vowels are more acoustically salient, and many researchers (Johnson, 2000; Kuhl, 1983; Werker & Tees, 1984, 1999) believe that vowel categories develop earlier than consonant categories. Thus it is possible that infants might perform very differently on our task than older learners. It may be that language development entails attending to information that is *more reliable* but *less acoustically salient*. The exact time course of any such age-based re-weighting of cues in representation remains an open question.

Conclusion

We have demonstrated clear similarities between lexical confusions in a relatively new artificial lexicon paradigm and lexical activation during on-line spoken word recognition paradigms. These similarities are somewhat surprising because our referent selection task is not typically viewed as a time-dependent task, whereas the eye movements taken to indicate lexical activation in fluent speech clearly tap on-line processing. We also have uncovered asymmetries in the relative weightings of segment types, and segment positions, in lexical identity. Given our results, it is interesting to consider the varying conditions under which lexical access occurs. Lexical access must be quite flexible if it is able to function, e.g., in high noise environments, in which case vowel information becomes predominant.¹² Nonetheless, it

¹² Creel, Aslin, and Tanenhaus (2005) taught listeners CVCV words with same-vowel and same-consonant competitors, either in noise or in the clear. Learners made more same-vowel confusions when learning and test were conducted in noise than when learning and test were noise-free, while overall error rates for both listening conditions were comparable. This suggested that listeners were weighting the more acoustically robust vowels more evenly with consonants in the noise condition.

seems clear that both early and mature lexical access are optimized to take advantage of the most reliable and fastest cues, though these vary across numerous situations. Now it is a matter of interest to uncover what situations take advantage of what cues, which will doubtless require an examination of the complex interplay between acoustic salience and acoustic differentiability in the formation of lexical representations.

Acknowledgments

This work was supported by NIDCD grant DC-005071 to M.K.T. and R.N.A., Packard grant 2001-17783 to R.N.A., and an NSF Graduate Research Fellowship to S.C.C.

Appendix A

Experiment 1

bamo	bami
tuzi	tuzer
gopa	gopu
seido	seida
niker	niku
pigu	piga
kobi	kobu
deisa	deisi
mutei	muto
zano	zaneı
beıka	gido
tibeı	nupeı
sogi	peizo
kanu	duma
maso	zoti
dipo	sumeı
geıma	pota
taki	zugeı
meıbo	nasu
bodi	kizu

Experiment 2

nado	gado
suka	buka
beızu	neızu
goti	poti
pimeı	simeı
tipu	mipu
zeısa	deısa
dogeı	togeı
kani	zani
mubo	kubo
bomu	zoba
daneı	padu
gipa	kizo
nugi	tukeı

Appendix A. (continued)

seito	meisi
teıso	geızi
pubi	duteı
sapeı	manu
bida	zimo
noku	koga

Experiment 3

CC		VV	
pona	puneı	nasi	tagi
dozeı	duzo	museı	pubeı
moki	meıka	tizo	mino
sobu	seıbi	bopu	kogu
gitu	geıto	keıta	seına
pino	peınu	padi	zapi
diza	dazi	gukeı	nuteı
mikeı	maku	diko	gibo
suba	sabo	sozu	domu
guti	gateı	zeıda	beıma

C		V	
zupo	zosi	dadeı	kanu
bopa	buma	busa	tudo
nodu	neısa	niba	simi
teıpi	tomeı	goza	pomo
kapu	kiso	deıpeı	neıni
zimu	zeıgu	gapo	saka
bida	bago	zuti	kubu
nami	nigeı	zigeı	bitu
tudeı	taseı	mozeı	toki
keıdo	kugi	meıgo	peısu

Experiment 4

CC		VV	
baʌsaıd ₃	basæd ₃	gasæd ₃	pabæd ₃
gupaıd ₃	gaʌpeıd ₃	tıbaı d ₃	pıgaıd ₃
pugađ ₃	pıgıd ₃	gobıd ₃	topıd ₃
sataıd ₃	saʌtıd ₃	bıgeıd ₃	pıteıd ₃
tubaıd ₃	tabeıd ₃	gupaʌđ ₃	pusaʌđ ₃
bısaıd ₃	bısođ ₃	batæđ ₃	sapađ ₃
gıpaıd ₃	gıpeđ ₃	gıtaıd ₃	bıpaıd ₃
pıgeıd ₃	pagađ ₃	bosıd ₃	sogıd ₃
sıteıd ₃	sutođ ₃	sıbeıd ₃	tıseıd ₃
tıbeđ ₃	taʌbod ₃	tugaʌđ ₃	sutaʌđ ₃

C		V	
beıpaʌđ ₃	bobıd ₃	taʌgad ₃	saʌgud ₃
gıtaıd ₃	gogıd ₃	paitad ₃	baıgod ₃
pıbıd ₃	pıpaʌđ ₃	geıbıd ₃	bæsaıd ₃
sıgud ₃	seıpıd ₃	peıgıd ₃	beıpud ₃
tıtıd ₃	tæıd ₃	bıtıd ₃	pıbıd ₃
baıtıd ₃	bogađ ₃	paʌsođ ₃	gaʌpıd ₃
gaisud ₃	geıbad ₃	saitud ₃	tasıd ₃

(continued on next page)

Appendix A. (continued)

C		V	
pɛtʊdʒ	peisʊdʒ	sɛbʌdʒ	tɛpʊdʒ
saisaʊdʒ	sɛbʌʊdʒ	teɪbʊdʒ	geisidʒ
taɪgʌdʒ	teɪpɪdʒ	ɡɪtʌdʒ	sɪpʊdʒ

Experiment 5

CC		VV	
oɪpʌn	uɪpeɪn	ʌnɪs	ʌtɪɡ
oʊdeɪz	uʊdoz	uɪmeɪs	uɪpeɪb
oɪmɪk	eɪmʌk	ɪtoz	ɪmʊn
oʊsʊb	eɪsɪb	oʊbʊp	oʊkʊɡ
ɪɡʊt	eɪɡʊt	eɪkʌt	eɪsʌn
ɪpʊn	eɪpʊn	ʌpɪd	ʌzɪp
ɪdʌz	ʌdɪz	uɡeɪk	uɪneɪt
ɪmeɪk	ʌmʊk	ɪdʊk	ɪɡʊb
ʊsʌb	ʌsʊb	oʊsʊz	oʊdʊm
ʊɡɪt	ʌɡeɪt	eɪzʌd	eɪbʌm

C		V	
ʊzʊp	oʊzɪs	ʌdeɪd	ʌkʊn
oʊbʌp	ʊbʌm	ʊbʌs	ʊtʊd
oʊnʊd	eɪnʌs	ɪnʌb	ɪsɪm
eɪtɪp	oʊtem	oʊɡʌz	oʊpʊm
ʌkʊp	ɪkʊp	eɪdeɪp	eɪnɪn
ɪzʊm	eɪzʊɡ	ʌɡʊp	ʌsʌk
ɪbʌd	ʌbʊɡ	ʊzɪt	ʊkʊb
ʌnɪm	ɪneɪɡ	ɪzeɪɡ	ɪbʊt
ʊteɪd	ʌteɪs	oʊmeɪz	oʊtɪk
eɪkʊd	ʊkɪɡ	eɪmʊɡ	eɪpʊs

Experiment 6

CC		VV	
oɪfʌn	uɪfeɪn	ʌnfɪs	ʌtɪɡ
oʊdfeɪz	uʊdfoz	uɪmfɛs	uɪpeɪb
oɪmfɪk	eɪmfʌk	ɪftʊz	ɪmfʊn
oʊsfʊb	eɪsfɪb	oʊbfʊp	oʊkfʊɡ
ɪɡfʊt	eɪɡfʊt	eɪkfʌt	eɪsfʌn
ɪpfʊn	eɪpfʊn	ʌpfɪd	ʌzɪp
ɪdfʌz	ʌdfɪz	uɡfeɪk	uɪneɪt
ɪmfɛk	ʌmfʊk	ɪdfʊk	ɪɡfʊb
ʊsfʌb	ʌsfʊb	oʊsfʊz	oʊdfʊm
ʊɡfɪt	ʌɡfeɪt	eɪzfʌd	eɪbfʌm

C		V	
ʊzfʊp	oʊzɪs	ʌdfeɪd	ʌkʊn
oʊbfʌp	ʊbfʌm	ʊbfʌs	ʊtfʊd
oʊnfʊd	eɪnfʌd	ɪnfʌb	ɪsfɪm
eɪtɪfɪp	oʊfeɪm	oʊɡfʌz	oʊpʊm
ʌkfʊp	ɪkfʊs	eɪdfeɪp	eɪnfɪn
ɪzfʊm	eɪzfʊɡ	ʌɡfʊp	ʌsfʌk
ɪbfʌd	ʌbfʊɡ	ʊzɪt	ʊkfʊb
ʌnfɪm	ɪnfɛɡ	ɪzfeɪɡ	ɪbfʊt
ʊtfeɪd	ʌtfeɪs	oʊmfeɪz	oʊtɪk
eɪkfʊd	ʊkfɪɡ	eɪmfʊɡ	eɪpfʊs

References

- Allopenna, P. D., Magnuson, J. S., & Tanenhaus, M. K. (1998). Tracking the time course of spoken word recognition using eye movements: evidence for continuous mapping models. *Journal of Memory and Language*, 38, 419–439.
- Bailey, T. M., & Hahn, U. (2005). Phoneme similarity and confusability. *Journal of Memory and Language*, 52, 339–362.
- Charles-Luce, J., & Luce, P. A. (1990). Similarity neighbourhoods 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 neighbourhoods in young children's receptive vocabularies. *Journal of Child Language*, 22, 727–735.
- Coady, J. A., & Aslin, R. N. (2003). Phonological neighbourhoods in the developing lexicon. *Journal of Child Language*, 30, 441–469.
- Connine, C. M., Titone, D., & Wang, J. (1993). Do the beginnings of spoken words have a special status in auditory word recognition? *Journal of Memory and Language* 32, 193–210.
- Content, A., Kearns, R. K., & Frauenfelder, U. H. (2001). Boundaries versus onsets in syllabic segmentation. *Journal of Memory and Language*, 45, 177–199.
- Creel, S. C., Aslin, R. N., & Tanenhaus, M. K. (2005). The salience of consonants and vowels in learning an artificial lexicon: the effects of noise. In *Proceedings of the ISCA workshop on plasticity in speech perception*.
- Cutler, A., Sebastián-Gallés, N., Soler-Vilageliu, O., & Van Ooijen, B. (2000). Constraints of vowels and consonants on lexical selection: cross-linguistic comparisons. *Memory & Cognition*, 28, 746–755.
- Dahan, D., Magnuson, J. S., & Tanenhaus, M. K. (2001). Time course of frequency effects in spoken-word recognition: evidence from eye movements. *Cognitive Psychology*, 42, 317–367.
- Dollaghan, C. A. (1994). Children's phonological neighborhoods: half empty or half full? *Journal of Child Language*, 21, 257–271.
- Dorman, M. F., Kewley-Port, D., Brady, S., & Turvey, M. T. (1977). Vowel recognition: inferences from studies of forward and backward masking. *Quarterly Journal of Experimental Psychology*, 29, 483–497.
- Dumay, N., Gaskell, M. G., & Feng, X. (2004). A day in the life of a spoken word. In K. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the twenty-sixth annual conference of the cognitive science society* (pp. 339–344). Mahwah, NJ: Erlbaum.
- Gaskell, M. G., & Dumay, N. (2003). Lexical competition and the acquisition of novel words. *Cognition*, 89, 105–132.
- Goldinger, S. D. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review*, 105, 251–279.
- Hillenbrand, J., Getty, L. A., Clark, M. J., & Wheeler, K. (1995). Acoustic characteristics of American English vowels. *Journal of the Acoustical Society of America*, 97, 3099–3111.
- Horii, Y., House, A. S., & Hughes, G. W. (1970). A masking noise with speech-envelope characteristics for studying intelligibility. *Journal of the Acoustical Society of America*, 49, 1849–1856.

- Johnson, C. E. (2000). Children's phoneme identification in reverberation and noise. *Journal of Speech, Language, and Hearing Research, 43*, 144–157.
- Jusczyk, P. W., Goodman, M. B., & Baumann, A. (1999). Nine-month-olds' attention to sound similarities in syllables. *Journal of Memory and Language, 40*, 62–82.
- Kuhl, P. K. (1983). Perception of auditory equivalence classes for speech in early infancy. *Infant Behavior & Development, 6*, 263–285.
- Ladefoged, P. (2004). *Vowels and consonants* (2nd ed.). Oxford: Blackwell.
- Lisker, L., & Abramson, A. S. (1964). A cross-language study of voicing in initial stops: acoustical measurements. *Word, 20*, 384–482.
- Luce, P. A., & Pisoni, D. B. (1998). Recognizing spoken words: the neighborhood activation model. *Ear and Hearing, 19*, 1–36.
- Macmillan, N. A., Goldberg, R. F., & Braidia, L. D. (1988). Resolution for speech sounds: basic sensitivity and context memory on vowel and consonant continua. *Journal of the Acoustical Society of America, 84*, 1262–1280.
- Maddieson, I. (1984). *Patterns of sounds*. New York: Cambridge.
- Magnuson, J. S., Tanenhaus, M. K., Aslin, R. N., & Dahan, D. (2003). The time course of spokeord learning and recognition: studies with artificial lexicons. *Journal of Experimental Psychology: General, 132*, 202–227.
- Marks, E. A., Moates, D. R., Bond, Z. S., & Stockmal, V. (2002). Word reconstruction and consonant features in English and Spanish. *Linguistics, 40*, 421–438.
- Marslen-Wilson, W. (1987). Functional parallelism in spoken word recognition. *Cognition, 25*, 71–102.
- Marslen-Wilson, W., & Zwitserlood, P. (1989). Accessing spoken words: the importance of word onsets. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 576–585.
- Marslen-Wilson, W., Moss, H. E., & van Halen, S. (1996). Perceptual distance and competition in lexical access. *Journal of Experimental Psychology: Human Perception and Performance, 22*, 1376–1392.
- McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18*, 1–86.
- Mehler, J., Dommergues, J. Y., Frauenfelder, U., & Segui, J. (1981). The syllable's role in speech segmentation. *Journal of Verbal Learning and Verbal Behavior, 20*, 298–305.
- Norris, D. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition, 52*, 189–234.
- Peterson, G. E., & Barney, H. L. (1952). Control methods used in a study of the vowels. *Journal of the Acoustical Society of America, 24*, 175–184.
- Redford, M. A., & Diehl, R. L. (1999). The relative perceptual distinctiveness of initial and final consonants in CVC syllables. *Journal of the Acoustical Society of America, 106*, 1555–1565.
- Samuel, A. G. (1991). Perceptual degradation due to signal alternation: implications for auditory pattern processing. *Journal of Experimental Psychology: Human Perception and Performance, 17*, 392–403.
- Singer, D'Oliveiro (2001). Available from: <http://www.princeton.edu/~bdsinger/old/SpeechSaver/>.
- Swingle, D., & Aslin, R. N. (2000). Spoken word recognition and lexical representation in very young children. *Cognition, 75*, 1–20.
- Swingle, D., & Aslin, R. N. (2002). Lexical neighborhoods and the word-form representations of 14-month-olds. *Psychological Review, 13*, 480–484.
- Swingle, D., Pinto, J. P., & Fernald, A. (1999). Continuous processing in word recognition at 24 months. *Cognition, 71*, 73–108.
- Van Ooijen, B. (1996). Vowel mutability and lexical selection in English: evidence from a word reconstruction task. *Memory and Cognition, 25*, 573–583.
- Werker, J. F., & Tees, R. C. (1984). Cross-language speech perception: evidence for perceptual reorganization during the first year of life. *Infant Behavior & Development, 7*, 49–63.
- Werker, J. F., & Tees, R. C. (1999). Experiential influences on infant speech processing: toward a new synthesis. *Annual Review of Psychology, 50*, 509–535.