

Modeling the word recognition data of Vitevitch and Luce (1998):

Is it ARTful?

Mark A. Pitt, Jay I. Myung, and Nicholas Altieri

Ohio State University

running head: ARTphone

word count=4028

Correspondence information:

Mark A. Pitt
Department of Psychology
1885 Neil Avenue
Ohio State University
Columbus, OH 43210-1222
office (614) 292- 4193
fax (614) 688-3984
pitt.2@osu.edu
<http://lpl.psy.ohio-state.edu>

Abstract

Vitevitch and Luce (1998) showed that the probability with which phonemes co-occur in the language (phonotactic probability) differentially affects the speed with which words and nonwords are named. Words with high phonotactic probabilities between phonemes were named more slowly than those with low probabilities. Just the opposite was found with nonwords. To reproduce this reversal in performance, a model would seem to require not only sublexical representations, but ones that are relatively independent of lexical representations, which is how ARTphone (Grossberg, Boardman, & Cohen, 1997) is designed. In this study, we used a technique called parameter space partitioning to analyze ARTphone's behavior and to learn not only whether it can mimic human behavior, but also to understand how. To perform best, differences in sublexical node probabilities must be amplified relative to lexical node probabilities to offset the additional source of inhibition (from top-down masking) that is found at the sublexical level.

Research in spoken word recognition focuses on delineating how the acoustics of a spoken word map onto a listener's memory representation of that word. One dimension on which models of recognition differ is whether smaller, sublexical representations, such as phonemes, biphones, or syllables are also formed during recognition. In models such as the dynamic cohort model (DCM: Gaskell & Marslen-Wilson, 2002), the neighborhood activation model (NAM: Luce & Pisoni, 1998), and the lexical access from spectra model (LAFS; Klatt, 1989), sublexical representations are absent, and considered to be unnecessary and potentially disruptive to successful recognition. In contrast, models such as TRACE (McClelland & Elman, 1986), Merge (Norris, McQueen, & Cutler, 2000), Shortlist (Norris, 1994), PARSYN (Luce, Goldinger, Auer & Vitevitch, 1998), and Adaptive Resonance Theory (ART; Grossberg, 1986) postulate at least one sublexical level of representation. Among other things, they help explain how nonwords and sub-word segments are perceived.

Evidence has been accumulating in support of the existence of sublexical representations, most recently in perceptual learning studies (Eisner & McQueen, 2005; Norris, McQueen, & Cutler, 2002). Participants in these experiments first perform a brief lexical decision task in which a subset of the words has the same final segment (e.g., /f/), but the segment has been replaced with one that is perceptually ambiguous (e.g., midway between /s/ and /f/). In a subsequent phoneme identification task using a /s-/f/ continuum, listeners display a bias in labeling ambiguous segments as /f/, suggesting their /f/ category was altered from the exposure in the initial lexical decision task. These data are parsimoniously accounted for by positing a sublexical representation, where perceptual learning is specific to a sublexical (e.g., phonemic) representation. How such segment-specific learning could be accounted for lexically (i.e., across

words) is not clear.

Additional support for sublexical representations comes from a series of studies by Vitevitch and Luce (1998, 1999, 2005). Listeners heard CVC words and nonwords that varied in phonotactic probability (the frequency with which segments co-occur in the language), which also correlates strongly with neighborhood density (the number of words in the language that are phonologically similar to a specific word). Across sets of stimuli and multiple tasks (e.g., naming, same-different judgment), facilitory effects of phonotactic probability were found when responding to nonwords, but inhibitory effects were found when responding to words. Nonword naming times were fastest when phonotactic probability was high versus low (see Lipinski & Gupta, 2005 for constraining conditions). Just the opposite was found with words, where naming times were slowest to words with high phonotactic probabilities between segments.

Vitevitch and Luce explain their findings by suggesting that the negative effect of phonotactic probability for words is due to lexical inhibition. Words in dense neighborhoods will receive much more inhibition than those in sparse neighborhoods. Nonwords, whether low or high in phonotactic probability, are likely to generate little lexical activity. Thus, any inhibition should be weak, enabling facilitory effects of phonotactic probability to emerge, which must be at a sublexical level because nonwords do not have lexical representations.

Vitevitch and Luce (1999) also note that the *reversal* of probability/density effects as a function of lexical status would seem to require that sublexical processing be relatively independent of lexical processing. Such independence is found in the ARTphone model (Grossberg, Boardman, & Cohen, 1997), and Vitevitch and Luce use the framework of ARTphone to explain how it could generate their results. Support for ARTphone's suitability,

and for the necessity of a sublexical representation, could be strengthened by demonstrating that the model can in fact reproduce the Vitevitch and Luce data. The purpose of the present study is to perform such a test of ARTphone, not just to determine whether it can mimic human behavior, but also to examine the extent to which the mechanisms hypothesized to be responsible for the empirical pattern (probability dominant sublexically and inhibition dominant lexically) operate in the way suggested by Vitevitch and Luce.

The ARTphone model

ARTphone (Grossberg et al, 1997) is an interactive activation model that has two main stages of processing. The first is *working memory*, and encodes the phonetic features of speech into phonemes. The second level represents chunks of speech of varying sizes, from phonemes to words (referred to generically as *list chunks*), that are stored in *short-term memory*. In contrast to models like TRACE and Merge, list chunks are not hierarchically arranged according to size to form sequentially ordered stages of processing. Rather, chunks of all sizes co-exist together in short-term memory, although they are differentiated by how they are wired. Items in working memory are connected to list chunks in short-term memory via bidirectional excitatory links, which is the source of the resonance that can lead to perception if sufficiently strong. Inhibitory links interconnect chunks of the same size. Larger chunks (e.g., words) mask (i.e., inhibit) smaller chunks (e.g., biphones and phonemes).

The network is depicted in Figure 1. Only the parts of short-term memory of most relevance for this investigation are shown in detail. There are four word nodes and eight biphone nodes. Not shown in detail are the seven phoneme nodes. For clarity, only the masking and a few of the lateral inhibition links are drawn, but the total number of such inhibitory links

(masking + lateral) impinging on each node is indicated by the numeral inside the node. The word nodes were interconnected to represent two levels of lexical density, low (zero neighbors; the word *bin*) and high (two neighbors, the words *can*, *ran*, *tan*). The corresponding biphones that make up the words are represented sublexically. Two additional sublexical chunks, *ba* and *at*, were included for the purpose of simulating nonword processing. Note that each word masks the two biphones that comprise it, and that *an* receives the most masking and *ba* and *at* none. Chunks of the same size inhibit each other if they overlap. For words, overlap was defined in terms of biphones. For biphones, it was defined in terms of phonemes.¹

Phonotactic probabilities were encoded in the network by multiplying the biphone and word bottom-up activation functions by a constant, whose value corresponded to the list chunk's probability in a corpus of English. Biphone probabilities were obtained from the Phonotactic Probability Calculator (Vitevitch & Luce, 2004). Following Vitevitch and Luce (1999), phonotactic probabilities for words were computed by averaging the two biphone probabilities that make up a word. Pilot simulations showed ARTphone to be highly sensitive to probability differences, so interval relations between item (word and biphone) probabilities could not be preserved, but ordinal relations were. The node with the highest probability was assigned a constant of 1.0 and the remaining nodes were scaled downward by a function of their phonotactic probability. Item probability was not held constant throughout the investigation. Rather, we varied the range over which the probabilities extended (close together or farther apart). The purpose of this procedure was two-fold: One was to identify an appropriate scaling of probability in which the empirical pattern was most robust; this exercise is necessary when introducing new factors into a model because it is not known beforehand the level of granularity

at which probability should be represented. The second reason was to learn how changes in probability affect model performance lexically and sublexically. The ranges and a description of their calculation can be found in the Appendix.

Vitevitch and Luce (1998) calculated neighborhood density by summing the log frequencies of a word's neighbors. To increase the similarity of the simulations to their experiment, bottom-up activation of word nodes was also modulated by word frequency in addition to phonotactic probability. The correlation between the two measures is very low ($r < .10$; Michael Vitevitch, personal communication), so information about word frequency is not redundant. Frequency counts were taken from Kucera and Francis (1967), and were converted to proportions using the same formula as with phonotactic probabilities (see the Appendix). Each value was multiplied by the word's phonotactic probability prior to being integrated with the bottom-up activation function for that word.

Input to ARTphone was a CVC, with the phonemes sequentially presented, each over three time units. To simulate coarticulation, the vowel and the final consonant overlapped their preceding phoneme by one time unit. The word inputs were *tan* and *bin*. Both have the same frequency but *tan* is from a higher density neighborhood (i.e., greater phonotactic probability) than *bin*. Although *ban* and *bat* are words, they served as nonword inputs for the simulations (note that they have no corresponding lexical nodes). These items differ in the probability of the second biphone, with *an* being higher than *at*. There is currently no mechanism in ARTphone for combining biphone chunks to yield a nonword node on the fly, so resonance functions of the second biphone (*an* and *at*) were used as the measure of nonword perception. The first biphone, *ba*, was held constant to ensure network behavior remained the same prior to presentation of the

final phoneme.

We first coded the network described in Grossberg et al (1997), validated the implementation by replicating the simulations in their paper (e.g., Figures 10 and 11), and then scaled up the model to the design described above. Parameters were held constant at the values listed in Figure 10 of Grossberg et al (1997) except that $\mu=2.2$ (it affects connection dynamics) and two parameters were added to the network, masking (described more fully below) and *kappa*, a constant that modulates the strength of top-down excitation (list chunk to input), which was set to 0.2 based on pilot simulations.

Analysis Methodology

We studied ARTphone's performance using a model analysis method called parameter space partitioning (PSP; Pitt, Kim, Navarro & Myung, 2006). In PSP, the range of data patterns a model can produce (e.g., the empirical pattern plus any others) in an experimental design is identified by repeating a simulation across the ranges of the model's parameters. Two of ARTphone's nine parameters were of interest, inhibition (between chunks of the same size) and masking (from large chunks to smaller chunks), so only they were varied. Pilot testing of various parameter values showed the model generated sensible patterns (e.g., functions with a single, identifiable resonance peak and no further resonance oscillations) when inhibition was restricted to the range 0-.15 and masking to the range 0-.30. Because only two parameters were varied, a simple grid search (Johansen & Palmeri, 2002) procedure was used to map the parameter space instead of the more sophisticated algorithm described in Pitt et al (2006).

To perform a PSP analysis, the results of Vitevitch and Luce (1998) must be defined quantitatively in terms of ARTphone performance. *Resonance* refers to the strength of the

interaction between phoneme inputs and a list chunk, with resonance increasing as the match between the two increases. Peak resonance is therefore a reasonable quantitative measure of the strength of evidence in favor of a particular list chunk, biphone or word. It is important to point out that predictions with behavioral measures such as reaction time correlate negatively with resonance strength: Longer reaction times (i.e., less efficient or slower processing) are equivalent to weaker resonance. To transform the results of Vitevitch and Luce into ARTphone performance, the predictions across word and nonword conditions must be reversed, so that faster reaction times correspond to greater resonance.

For ARTphone to reproduce the Vitevitch and Luce data, low density words (*bin*) should achieve a higher peak resonance than high density words (*tan*), and high probability nonwords (*ban*) should achieve higher resonance than low probability nonwords (*bat*). This empirical data pattern is one of nine possible in the 2x2 design of Vitevitch and Luce (1998) if equalities and inequalities between conditions are considered (node resonances were classified as equal if their peak activations differed by less than .02 units). All nine patterns are shown in the Table in Figure 2. To qualify as one of these patterns, all four inputs had to exceed a minimum level of activation for a given set of parameter values. For word nodes, this value was 0.2. For biphones, it was 0.1. Simulations in which activation failed to reach this level were labeled as invalid (pattern 10). Note that pattern 3 is the empirical pattern, and pattern 7 is the opposite pattern. Matlab code for the simulations and analyses is available from the authors. Included with these files is a document that describes the model in more detail, including equations.

A PSP analysis is ideal for studying model behavior because it can reveal how many of the nine patterns the model can produce. The analysis also provides information about the

centrality of each data pattern in the model, in the form of the size of the area occupied by that pattern in the model's parameter space. Data patterns that occupy a large region are more representative of model behavior than those that occupy a small region. This analysis assumes a uniform distribution of prior probabilities of parameter values in the specified ranges.

Results and Discussion

The goal of this investigation was to learn not just whether ARTphone can generate the empirical pattern, but how. Vitevitch and Luce (1999) suggest that it arises because probability dominates sublexically and inhibition dominates lexically. Inspection of the model's design in Figure 1 suggests this data pattern might be difficult to produce. Because there are more direct sources of inhibition sublexically (lateral and masking) than lexically (lateral only), one might think that inhibition should dominate sublexically. However, this prediction ignores potential influences of phonotactic probability, which could reverse this built-in bias and thus generate the empirical pattern. Can it? To answer this question, as well as examine how these multiple forces interact in the model, we ran the PSP analysis several times while varying the range of phonotactic probabilities.

In the first analysis, all sublexical and lexical nodes were fixed at the same phonotactic probability and frequency (1.0, the maximum). If differences in phonotactic probability are necessary to produce the empirical pattern, then the empirical pattern should not be found in the analysis, which is indeed the case. Shown in panel a of Figure 2 is a plot of the regions in ARTphone's parameter space (i.e., each combination of inhibition and masking settings) occupied by each of the nine possible data patterns in the experimental design. The lines separating the regions had to be hand-edited, so their positions should be considered

approximate.

Without differences in probability, only masking and inhibition drive processing. Sublexically, this means that a node with fewer inhibitory links (*at*) achieves greater resonance than a node with more links (*an*), exactly opposite of what is found with humans. This result (patterns 1, 4, 7) occupies most of the parameter space. Only at very low levels of masking does *an* achieve enough resonance to be considered equivalent to *at* (patterns 2, 5, 8). For the lexical items (*tan* and *bin*), all three outcomes (High<Low, High=Low, High >Low) are possible depending on the combination of inhibition *plus* masking. Under some parameter settings, the effects of masking can percolate back up to alter lexical node activation.

In the next two analyses, probability differences between nodes (sublexical and lexical) were introduced by a small amount (panel b) and a larger amount (panel c; see Appendix for ranges). In panel b, the empirical pattern is present in the upper left corner. In this rather small region of parameter space, differences in probability at the sublexical level are large enough to offset weak masking effects (combined with strong inhibitory effects) and yield greater resonance for high than low probability chunks. At the same time, at the lexical level, inhibition is great enough in this region to achieve just the reverse effect: Inhibition in a high-density neighborhood overwhelms any advantage of greater phonotactic probability so that the word in the low density neighborhood achieves the highest resonance. Nevertheless, the empirical data pattern is rather fragile with these probability settings, disappearing quickly if masking is increased. As in panel a, ARTphone generates the other lexical patterns (High =Low, High >Low) depending on inhibition and masking level.

Further comparison of panels a and b shows ARTphone to be extremely sensitive to the

phonotactic probability of the words. After introducing small differences in probability, patterns 7-9 occupy even larger portions of the parameter space. The higher density word (*tan*) easily overcomes the inhibition from its somewhat small neighborhood, leading to an increase in area occupied by the High>Low lexical pattern.

In panel c, where the range of probabilities was greater, three patterns dominate (7, 8, 9). Relative to panel b, what has happened is that these patterns occupy even more of the parameter space, further extending their reach upward into regions of greater inhibition. Again, when phonotactic probability differences are introduced, the higher density word quickly dominates, which is why only the High > Low pattern is found. Sublexically, as masking decreases, the pattern again shifts from High < Low (7) to High > Low (9), re-affirming the need for small amounts of masking to generate the empirical nonword pattern.

Two observations emerge from these first three analyses. One is that ARTphone is highly sensitive to differences in phonotactic probability. Perhaps because of this, simply introducing them is not enough to generate the empirical pattern robustly. The other is that differences in phonotactic probability appear much more potent lexically than sublexically. One reason for this was alluded to earlier: There is only one type of inhibition lexically (lateral), but two sublexically (lateral and masking). ARTphone's rather weak ability to generate the empirical pattern thus far could be a result of failing to compensate adequately for the strength of this difference.

One way to offset the difference is to vary phonotactic probability more widely across sublexical than lexical items. Such a configuration of the model is in keeping with what Vitevitch and Luce (1999) suggest is necessary to generate the empirical pattern, where

inhibition should dominate lexically (because phonotactic effects are weak) and probability should be strongest sublexically (to offset two sources of inhibition).

Accordingly, a final PSP analysis was run with sublexical probability varying much more widely than lexical. Panel d contains the results. This change had the intended effect of substantially increasing region 3, the area occupied by the empirical pattern. Additional analyses using other combinations of probability ranges did not increase this area substantially. The graph on the right (panel e) shows ARTphone simulating the empirical pattern using masking and inhibition values from the center of region 3 (masking=.04, inhibition=.13) .

Although it is re-assuring that ARTphone can generate the empirical pattern, it is somewhat disconcerting that this configuration of ARTphone can also produce all of the other data patterns in the experimental design expect pattern 1. In particular, that pattern 7, which is opposite of the empirical pattern, is still the most stable suggests that the proclivities of ARTphone have to be counteracted to perform like listeners. Together, these findings can cast doubt on ARTphone's suitability.

This conclusion holds only if the parameterization of the model remains as is, with much of it unused. At the start of this study, we purposefully chose ranges of the inhibition and masking parameters that generated reasonable resonance functions (e.g., panel e). The knowledge gained from the current simulations suggests that the range of both parameters could be restricted even further, perhaps keeping masking below .15 and inhibition above .06. This would have the benefit of decreasing the number of data patterns produced as well as increasing the area of the empirical pattern. Such a change is particularly justifiable if ARTphone uses parameter values in this range when simulating other data sets.

ARTphone's ability to produce a wide range of data patterns is not unexpected given its architecture. The virtual independence of lexical and sublexical chunks, in that they are connected only via one-way masking links, gives the network flexibility in producing a range of data patterns. We (Pitt et al, 2006) have observed similar flexibility in the Merge model of phoneme perception (Norris et al, 2000), which shares structural similarities with ARTphone. Phoneme input is fed simultaneously to a lexical stage and a phoneme decision stage in Merge. Excitatory links connect the lexical stage to the decision stage. As with ARTphone, the independence of these two higher-level stages enables the model to produce a range of data patterns. A downside of this flexibility, which is visible in Figure 2, is that less of the parameter space will be occupied by the empirical pattern.

In the introduction, we mentioned that the independence of the sublexical and lexical list chunks may be a reason why ARTphone would be able simulate the Vitevitch and Luce (1998) data. The results of this study suggest that masking performs a key role in simulating the empirical pattern. It must be present, and only at low levels. Between-level masking is a unique feature of ARTphone that distinguishes it from other models. For example, in a model such as TRACE, connections between lexical and sublexical levels are excitatory in both directions, making processing of shared phonemes mutually reinforcing. Unless TRACE can mimic masking in another way, the Vitevitch and Luce data could pose a challenge to simulate.

One issue that should be the focus of future work is the fact that the empirical pattern was found only when there were minimal differences in phonotactic probability across words. In fact, the area occupied by the empirical pattern actually increases slightly (expanding into region 6) when the final analysis (panel d) is rerun with word probabilities equated. Ideally this should not

happen. ARTphone's small lexicon might be to blame. In a larger network, more and longer words would be present. Longer list chunks would mask smaller ones. This additional source of inhibition could be offset by increasing probability differences between words.

The results of Vitevitch and Luce (1998) have been around for almost a decade, yet we know of no model that has simulated this challenging pattern of data. ARTphone can do so, and the means by which it succeeds nicely correspond to what appears to drive the effect in listeners. Finally, the PSP analyses bring to light the inner workings of ARTphone so that we understand how.

References

- Eisner, F., & McQueen, J.M. (2005). The specificity of perceptual learning in speech processing. Perception & Psychophysics, 67, 224-238.
- Gaskell, M.G., & Marslen-Wilson, W.D (2002). Representation and competition in the perception of spoken words. Cognitive Psychology, 45, 220-266.
- Grossberg, S. (1986). The adaptive self-organization of serial order in behavior: Speech, language, and motor control. In E. C. Schwab & H. C. Nusbaum (Eds.), Pattern recognition by humans and machines: Vol. 1. Speech perception, New York: Academic Press.
- Grossberg, S., Boardman, I., Cohen, M. (1997). Neural dynamics of variable-rate speech categorization. Journal of Experimental Psychology: Human Perception and Performance, 23, 483-503.
- Johansen, M. A. & Palmeri, T. J. (2002). Are there representational shifts during category learning? Cognitive Psychology, 45, 482-553.
- Klatt, D.H. (1980). Speech perception: A model of acoustic-phonetic analysis and lexical access. In R.A. Cole (Ed.), Perception and production of fluent speech. Hillsdale, NJ: Erlbaum.
- Kucera, H., & Francis, W. N. (1967). Computational analysis of present-day American English. Providence, RI: Brown University Press.
- Luce, P.A., Goldinger, S.D., Auer, Jr., E.T., & Vitevitch, M.S. (1998). Phonetic priming, neighborhood activation, and PARSYN. Perception & Psychophysics, 62, 615-625 .
- Luce, P.A., & Pisoni, D.B. (1998). Recognizing spoken words: The neighborhood activation model. Ear and Hearing, 19, 1-36.

- McClelland, J.L., & Elman, J.L. (1986). The TRACE model of speech perception. Cognitive Psychology, 18, 1-86.
- Norris, D. (1994). SHORTLIST: A connectionist model of continuous speech recognition. Cognition, 52, 189-234.
- Norris, D., McQueen, J.M., & Cutler, A. (2000). Merging information in speech recognition: Feedback is never necessary. Behavioral & Brain Sciences 23, 299-370).
- Norris, D., McQueen, J.M., & Cutler, A. (2002). Perceptual learning in speech. Cognitive Psychology, 47, 204-238.
- Pitt, M.A., Kim, W., Navarro, D., & Myung, I. (2006). Global model analysis by parameter space partitioning. Psychological Review, 113, 57-83.
- Vitevitch, M.S., & Luce, P.A. (1998). When words compete: Levels of processing in spoken word perception. Psychological Science, 9, 325-329.
- Vitevitch, M.S., & Luce, P.A. (1999). Probabilistic phonotactics and neighborhood activation in spoken word recognition. Journal of Memory & Language, 40, 374-408.
- Vitevitch, M.S., & Luce, P.A. (2004). A Web-based interface to calculate phonotactic probability for words and nonwords in English. Behavioral Research Methods, Instruments, & Computers, 36, 481-487.
- Vitevitch, M.S., & Luce, P.A. (2005). Increases in phonotactic probability facilitate spoken nonword repetition. Journal of Memory & Language, 52, 193-204.

Appendix

Values used to represent biphone and word probabilities in the simulations, which modulated the bottom-up flow of activation (i.e., upward-pointing arrows in Figure 1). Rather than choose sets of values arbitrarily, they were generated using the equation $a \cdot \log(prob) + 1 - a \cdot \log(b)$, where a was the slope of the logarithmic function, $prob$ was the phonotactic probability of the biphone or word, and b was set to the highest value in each category (0.0144 for biphones and 0.0133 for words) to ensure the intercept was equal to 1. Biphone probabilities were taken from the Phonotactic Probability Calculator (Vitevitch & Luce, 2004). Following Vitevitch and Luce (1999), phonotactic probabilities for words were computed from the same source by averaging the two biphone probabilities of each word.² Word frequency values (in parentheses) were rescaled from the values in Kucera and Francis (1967), and then plugged into the above equation with $a=.05$ and $b=.01772$.

Biphones	Phonotactic Probability	Simulation Probability Values			
		fixed a=0	small range a=.08	larger range a=.15	Sublexical > Lexical a=.3
an	0.0144	1	1	1	1
ca	0.0144	1	0.9942	0.9892	0.9784
in	0.0095	1	0.9855	0.9729	0.9458
ba	0.0059	1	0.9690	0.9419	0.8837
at	0.0059	1	0.9690	0.9419	0.8837
ra	0.0050	1	0.9632	0.9311	0.8622
bi	0.0041	1	0.9564	0.9182	0.8363
ta	0.0039	1	0.9546	0.9149	0.8298
Words		a=0	a=.23	a=.60	a=.017
can (1.00)	0.0133	1	1	1	1
ran (0.943932)	0.0097	1	0.9685	0.9178	0.9977
tan (0.885289)	0.0092	1	0.9626	0.9025	0.9972
bin (0.885289)	0.0068	1	0.9330	0.8252	0.9950

¹ The pragmatics of trying to create a small but realistic network in which neighborhood density is manipulated prompted us to define overlap differently lexically and sublexically. The connectivity of this network would not have to change if overlap were defined solely in terms of phonemes. Instead, different words could be substituted. Also, although a much larger network could yield more realistic results, all links in ARTphone must be wired by hand, which makes it tricky to scale up the model further. Over 100 data vectors were required for this implementation.

² Readers familiar with Vitevitch and Luce (1999; Experiment 1) might notice that the word probabilities for *tan* and *bin* (0.0092 vs. 0.0068) are more similar to one another than the mean biphone probabilities from the high- and low-probability conditions in their experiment, which were 0.0143 and 0.0006, respectively. Had we used these means instead, only the values of *a* in the Appendix (for words only) would have changed, not the simulation results. Recall that multiple ranges of probability values were used for the purpose of determining the granularity at which word and biphone probability itself should be represented in ARTphone to generate the empirical pattern. When these means are substituted (and appropriate probabilities are assigned to *can* (0.015) and *ran* (0.0147)), a graph very similar to panel d is produced with $a=0.00347$ and $b=0.015$.

Author Notes

Mark A. Pitt, Jay I. Myung, Department of Psychology, Ohio State University. Nicholas Altieri, Department of Psychology, Indiana University.

This work was supported by research grant R01-MH57472 from the National Institute of Mental Health, National Institute of Health. NA was also supported by an Undergraduate Research Scholarship from OSU. Matlab code and documentation of the model can be found at <http://lpl.psy.ohio-state.edu/software.html>.

Correspondence concerning this article should be addressed to Mark Pitt, Department of Psychology, Ohio State University, 1885 Neil Avenue, Columbus, OH, 43210. Electronic mail should be sent to Pitt.2@osu.edu

Figure Captions

Figure 1. Schematic illustration of ARTphone with details of the connectivity in short-term memory. The most relevant inhibition and masking links between nodes are shown. Numerals inside the nodes denote the number of impinging inhibitory links.

Figure 2. The Table contains the nine data patterns that are possible in the 2x2 design of Vitevitch and Luce (1998). The relationship between high and low probability nonwords (sublexical items) is represented across columns. The relationship between high and low probability words is represented across rows. Each pattern is identified with a numeral, which is referenced in the text. Regions labeled 10 indicate an invalid data pattern. Starting with pattern 7 (opposite of the empirical pattern), as you move upward and leftward in the Table, the patterns become more human-like (pattern 3). Panels a-d show the range and extent of the data patterns that ARTphone generates when simulations are run over the specified range of inhibition and masking settings. These four graphs differ in the range of nonword and word probabilities. Panel e contains output from ARTphone using parameter and probability settings that generate the empirical pattern.

Figure 1

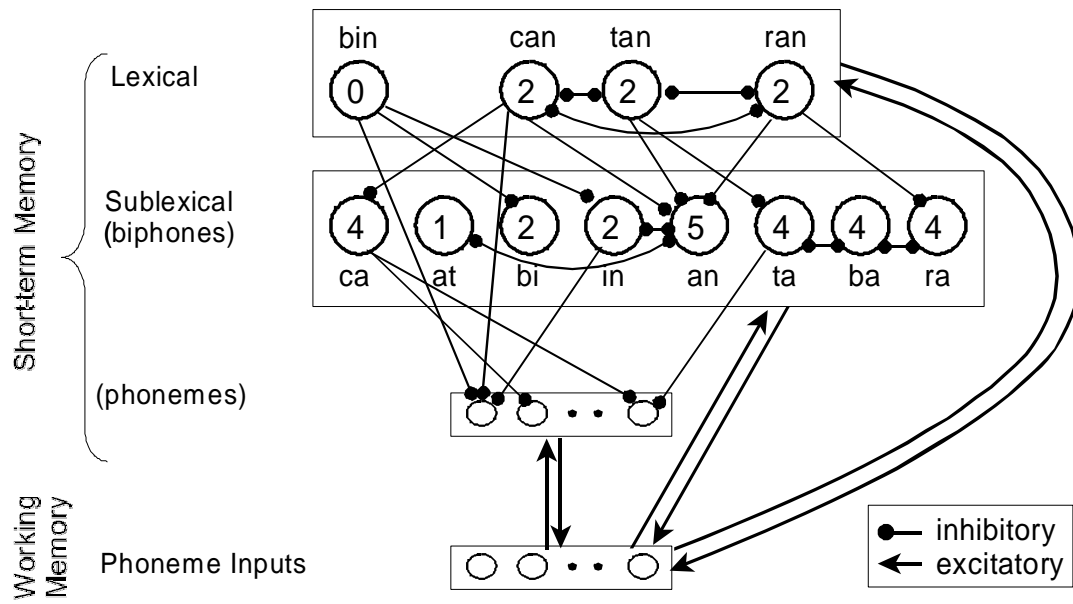


Figure 2

