

Reading Aloud Multisyllabic Words: A Single-Route Connectionist Model for Greek

Konstantinos D. Outos (outosk@hol.gr)

Graduate Program in Cognitive Science, Athens University Campus
GR-15771 Athens, Greece

Athanassios Protopapas (protopap@ilsp.gr)

Institute for Language & Speech Processing, Artemidos 6 & Epidavrou
GR-15125 Maroussi, Greece

Abstract

Multisyllabic word reading has received little attention in existing computational models, which are designed for English. The Greek language uses mainly multisyllabic words, while its orthography is feedforward consistent and includes a stress diacritic. We present here a computational model of reading aloud Greek words and nonwords, based on the connectionist “triangle” model, adapted to the Greek orthography, using a novel input representation. The network displays several effects from the word reading literature and successfully assigns stress. Further investigations are underway.

Keywords: Multisyllabic words; stress assignment; triangle model; reading aloud; nonword reading; connectionist; Greek.

Introduction

Computational Models of Reading Aloud

It has been 20 years since Seidenberg & McClelland (1989) presented the first “triangle” connectionist model of reading aloud. Since then, a variety of computational models have been presented, single-route (Plaut, McClelland, Seidenberg & Patterson, 1996; Harm & Seidenberg, 2001, 2004) or dual-route (DRC: Coltheart, Rastle, Perry, Langdon & Ziegler, 2001; CDP+: Perry, Ziegler, & Zorzi, 2007; Zorzi, Houghton, & Butterworth, 1998). These models differ in many respects, such as the existence of pre-defined grapho-phonemic decoding rules versus learning procedures, and the localist or distributed nature of lexical representations.

The triangle family of models has the advantage of discovering, during a training phase, regularities in the relations between orthographic and phonological representations in a set of words. The regularities are then generalized to novel stimuli, such as nonwords. Three implementations of this approach have been reported (Seidenberg & McClelland, 1989; Plaut et al., 1996; Harm & Seidenberg, 2001, 2004), using different representations, leading to differences in reading performance (Seidenberg & Plaut, 2006).

Most existing models have focused on reading English monosyllabic words. Problems associated with multisyllabic words, such as syllabification and stress assignment, have led to their exclusion. Likewise, most available data on human reading concern monosyllabic English words, facilitating comparative model assessment. To present monosyllabic words on a connectionist input layer, Harm & Seiden-

berg (2001, 2004) used slots corresponding to single letters separated into subsyllabic units. Onset and coda consonants were placed around the central vowel nucleus. This representation aimed to minimize the dispersion problem (Plaut et al., 1996), while allowing the model to “capture the fact that phonemes in different positions sometimes differ phonetically” (Harm & Seidenberg, 1999, p. 493).

Exclusive attention to monosyllables, and corresponding syllable-based representations, pose severe constraints to extension towards multisyllabic words. Specifically, to retain the existing structure, words must be pre-syllabified before being presented to the model. However, this limits the applicability of this approach to situations in which orthographic syllabification is possible.

The Greek Orthography

The Greek orthography is feedforward consistent and predictable to a large extent (about 95%; Protopapas & Vlahou, in press). Most graphemes can be mapped unambiguously to single phonemes when context is taken into account. The only substantial source of inconsistency concerns words containing the CiV pattern, that is, an unstressed grapheme normally mapping to /i/ when it follows a consonant and precedes a vowel (Protopapas & Vlahou, in press). In such cases there are two possible pronunciations, one of which contains an /i/ and another which contains a palatal consonant. The correct pronunciation is lexically determined. In rare cases, this situation leads to homographs. For example, the Greek words for “permission” and “empty” are both written as άδετα. However, “permission” is the three-syllable word /a.ði.a/ with an [i] forming the nucleus of the second syllable, whereas “empty” is the two-syllable word /a.ðja/ with the palatal consonant [j]. Therefore, the CiV phenomenon affects not only grapho-phonemic consistency but also orthographic syllabification as well.

In Greek, lexical stress always falls on one of the last three syllables and is affected by morphology (Revithiadou, 1999). Stress is orthographically marked with a special diacritic on every word with two or more syllables (Petrounias, 2002). This diacritic also disambiguates certain vowel digraphs, therefore it is necessary to include in orthographic representations, along with diaeresis.

These characteristics make the vowel-centered syllabic slot representation unsuitable for Greek multisyllabic words,

Whereas the homograph /'a. ðja / is represented as

Phonological:

___ a ___ | ___ ð j a ___ | ___ | ___ | ___ | ___ | ___
cccvc|cccvc|cccvc|cccvc|cccvc

Stress:

a - - - -

Training the Network

The training corpus was a list of 120,745 word types with 2–5 syllables, 4–10 letters long, from an online written text corpus of about 30 million words (Hellenic National Corpus; hnc.ilsp.gr). The corresponding pronunciations were derived using a grapheme-to-phoneme transcription model developed for text-to-speech applications (Chalamandaris, Raptis, & Tsiakoulis, 2005). Words were presented to the model proportionally to their frequency of appearance in that corpus after logarithmic compression and conversion to relative probability, following Harm & Seidenberg (1999, Equation 1, p. 495), to allow low frequency words to appear during training. The relative probability of all words with less than .3 occurrences per million was set to 0.2.

Beginning with random initial weights, 9,000,000 training trials were given. On each trial, a word was randomly chosen from the training list and was presented to the model if a random number from a flat [0,1] distribution did not exceed its transformed frequency. The orthographic representation of the word was fixed on the input layer for 10 “ticks” (i.e., activation update cycles through the network). The phonological output after 12 ticks was compared to the target output. Connection weights were adjusted using continuous recurrent backpropagation, with a learning rate equal to 0.1.

A subset (about 10%) of the training corpus (12,017 words) was retained and used to track performance during training every 50,000 trials (plotted in Figure 2). Each word was presented to the network for 18 ticks. A response was considered correct if produced within that time. More ticks were allowed in testing than in training, because longer words might require longer times for the network to settle to the correct pronunciation. Pilot trials showed significant improvement using 18 ticks for testing, compared to 12 ticks. Further increases (up to 45 ticks) did not improve performance significantly. Only 12 ticks were used during training, because this led to faster training, presumably due to increased pressure for the network to learn.

Post-training Tests

To study the effects of the most important variables known to affect word reading times in human participants (Balota, Yap, & Cortese, 2006), a set of 150 words were selected to span a range of length, frequency, neighbourhood size, and bigram probability. The words were chosen so that the intercorrelations between these basic variables were minimal (Spearman’s $\rho < 0.2$) and not statistically significant, in order to isolate individual effects. A corresponding set of 150 nonwords were constructed, based on these words, with similar properties and variable ranges, taking care to avoid

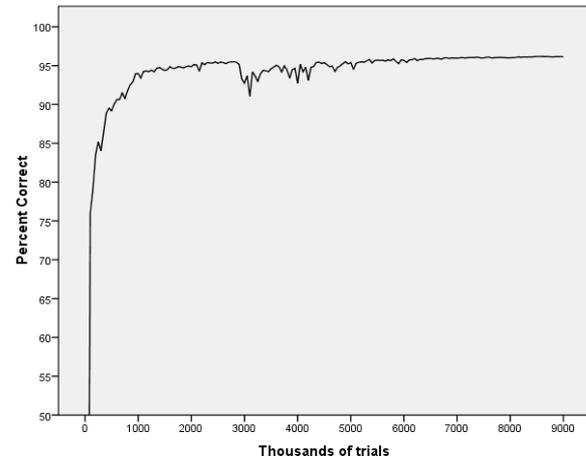


Figure 2. Word reading accuracy during training (%).

resemblance of nonwords to specific words. Human reading performance data for these stimuli are currently being collected, for future comparison with model performance.

After training, these words and nonwords were presented to the model with and without a stress diacritic, to assess its performance on segmental pronunciation and on stress assignment. Response time was measured, for correct responses only, by the number of ticks it took the model to achieve a representation of the correct response. The orthographic input was presented to the input layer for 18 ticks.

The variables onto which response time were regressed included: probability of appearance in training, as a measure of frequency (words only); length (in letters); number of syllables, number of orthographic neighbours (words only), cumulative bigram (letter) probability, mean orthographic syllable frequency, and orthographic transparency measured as minimum non-directional grapheme-phoneme type probability (least transparent sonograph; Spencer, 2007). Of these, only syllabic frequency differed significantly between words and nonwords ($t(298) = 2.007, p = .01$).

Results

Overall, 96.14% of words were read correctly after 9m trials. Performance may increase with further training, as no asymptote is apparently reached (Figure 2). 2.37% of the errors were stress assignment errors, that is, stress produced on a vowel other than the one marked with the diacritic (e.g., /ði.i.'li.ze/ instead of /ði.i.li.ze/). 4.53% were both segmental and stress assignment errors in that stress was erroneously produced on an incorrect vowel. The rest were segmental errors, that is, incorrect phonemes. When a stressed vowel was produced incorrectly, it was nevertheless stressed, showing that stress diacritic position was interpreted correctly (e.g., /ðja.vo.'lis/ instead of /ðja.vo.'lɛs/). 18.97% of the errors concerned stress assignment combined with a segment displacement, usually resulting from incorrect parsing of a CiV. In such cases, the stress diacritic was placed on the correct vowel at its new position, showing that

the model can properly combine letter and diacritic information (/fi.'ε.stes/ instead of /fçε.stes/).

To examine the effects of the aforementioned predictor variables, simple regressions were employed. This allows reliable detection of the most important effects, because the main variables were specifically uncorrelated in the test sets. Multiple regressions will be used once the critical variables are identified in the analysis of human data.

Of the 150 test words, presented with a stress diacritic, 98% were read correctly, with an average response time (RT) of 4.44 ticks. Linear regression, separately for each predictor, showed that number of letters accounted for 4.4% of RT variance ($R^2 = .044$, standardized $\beta = .084$, $p = .011$), and bigram probability for 4.5% ($\beta = -.028$, $p = .010$).

Of the 150 nonwords presented with a stress diacritic, 92% were read correctly, at a mean RT of 4.54 ticks. Number of letters accounted for 5.6% of RT variance ($\beta = .116$, $p = .005$), number of syllables for 3.2% ($\beta = .179$, $p = .035$), and bigram probability for 7.5% ($\beta = -.046$, $p = .001$).

These results are summarized in the following table.

Table 1: Summary of simple regression results for items presented with a stress diacritic. Significant predictors are presented in order of variance proportion accounted for.

Predictors	Words	Nonwords
Significant	Bigram prob. N Letters	Bigram prob. N Letters N Syllables
Not significant	N Syllables Frequency Syllable frequency N Neighbors Transparency	Syllable frequency Transparency

The RT difference between correct readings of words and nonwords was not significant ($t(283) = -1.457$, $p = .148$).

Of the 150 words presented without a stress diacritic, 53.33% were read correctly with a mean RT of 5.11 ticks, significantly slower than when presented with a stress diacritic ($t(80) = -3.267$, $p = .002$). Sixty errors were segmentally correct but incorrectly stressed words, including 6 responses with no stressed letter and 54 stressed at a different position. For the words produced correctly, the model assigned stress 27.5% on the final syllable, 42.5% on the penult, and 30% on the antepenult, a uniform distribution of stress ($\chi^2(2) = 3.100$, $p = .212$). Taking into account all words, stress assignment to the final, penult, and antepenult was 30.7%, 44.5%, and 24.8%, respectively (significantly nonuniform, $\chi^2(2) = 8.423$, $p = .015$). The correct stress positions for all words in the testing set were 30.7% on the final, 40% on the penult, and 29.3% on the antepenult; and for words stressed incorrectly by the model, 35.1%, 47.4%, and 17.5%, respectively. The corresponding proportions for the entire corpus, considering multisyllabic words only, are 30%, 44.9%, and 25%, respectively (Protopapas, 2006). Two words, which formed a CiV pattern when the stress

diacritic was removed, were read by the model with the alternate pronunciation of the CiV and were stressed appropriately considering the segment pattern produced.

Of the 150 nonwords presented without a stress diacritic, 42.67% were read correctly, with a mean RT of 5.3 ticks, significantly slower than when presented with a stress diacritic ($t(63) = -2.019$, $p = .048$). The model stressed nonwords 28.7% on the final syllable, 44.1% on the penult and 27.2% on the antepenult, a nonuniform distribution ($\chi^2(2) = 7.162$, $p = .028$).

Discussion

This is a first attempt toward a computational model of reading aloud Greek words and nonwords. The number of monosyllables in Greek is very small: fewer than 500 types were reported by Protopapas & Vlahou (in press), most of which were unrepresentative of the language in being either function words or recent loans. Therefore extension of existing approaches to multisyllabic representations was necessary in order to capture the major characteristics of this language. The Greek orthography is highly consistent for reading, with the exception of the CiV pattern. On the one hand, the high consistency makes the task of mapping letter sequences to phoneme sequences easier. On the other hand, the presence of the CiV phenomenon dictated a substantial change to the design of the model's representations, because pre-syllabification is not possible. The new design has the additional benefit that the number of syllables is not limited, as long as computational resources can handle the training. The concomitant drawback is susceptibility to the dispersion problem, because the same graphophonemic mappings must be learned repeatedly in different input-output slot positions.

Our novel orthographic representation bears some interesting features. The lack of pre-syllabification forces the model to learn mappings that might otherwise be distinguished by their subsyllabic position. The model must learn to map the letter sequences at the input layer to the syllabified output at the phonological output layer. This is not trivial, because letter positions are not fixed at the input, as they depend on word length and morphology. It is especially complex for successive vowel letters (up to 6 slots for single graphemes or digraphs) mapping to multiple syllables. However, even though the model is forced to learn the same mappings at several different slots, this does not seem to pose a serious problem for word reading performance or generalization to nonwords. This may be due to the relatively simple grapheme-phoneme mappings of the Greek orthography. It remains to be investigated whether the model learns to read in the same way as Greek readers do.

Figure 2 shows that early in training (100,000 trials) the model can already read correctly a considerable proportion of words (about 76%). The number of word types presented to the model (120,745), relative to the number of training trials (9 million), is huge, compared to the 3,123 words and 1 million trials in Harm & Seidenberg (1999), 6,103 words and 1.5 million trials in Harm & Seidenberg (2004), 5,870 words and 1 million trials in Zevin & Seidenberg (2006),

and 9,911 words and 1.2 million trials in Pagliuca & Monaghan (in press). Despite the low ratio of tokens to types, the lack of syllabification, and the dispersion over slots, the model can read correctly more than 96% of the training corpus. This may be due to high orthographic transparency.

Due to the frequency-modulated random selection procedure, many words that are read correctly were never presented to the model during training, so they must be read by grapheme-to-phoneme conversion, as nonwords. It is instructive to examine the model's response to various letter strings containing the CiV pattern, because there is no rule for the CiV, either in terms of a statistical preponderance (Protopapas & Vlahou, in press) or in human participants' reading behavior (Protopapas & Nomikou, 2009). The results of the training test show instances of words read with the incorrect alternative pronunciation (but not with unrelated phonological outputs). These cases concerned low-frequency words that were never or little presented to the model during the training, so they are functional nonwords. Further tests, with controlled sets of words, are underway.

Another novel feature of this model is stress marking in the orthographic representation, corresponding to the diacritic of Greek orthography. Although this is not the first attempt to model stress assignment in reading (e.g., Monaghan, Arciuli & Seva, 2008; Rastle & Coltheart, 2000) or to include an orthographic representation of stress (Pagliuca & Monaghan, 2009, in press), it is probably the first attempt to consider a distinct orthographic representation for the stress diacritic itself while retaining vowel letter identity. The model seems to have learned the constraints on stressed syllables. No stress assignment error was observed on a syllable earlier than the antepenultimate even though syllable positions were not fixed or right-aligned. Analysis of incorrectly stressed words showed that the model has a stress "preference" for the penultimate syllable, like humans, in both words and nonwords. There was no significant distortion of stress assignment toward any syllable, indicating that the model does not assign stress randomly but follows the distribution of stress positions seen on Greek words.

The network seems to have learned to make a connection between stressed vowels and the stress diacritic, even when the stressed vowel was incorrectly produced or placed in the wrong syllable. In such cases, stress followed the vowel, either by changing vowel or by changing position. Only 0.1% of the total training corpus was read with the correct segmental pronunciation and incorrect stress. This means that the model has learned to use the stress diacritic. On the other hand, the model's stress assignment performance deteriorated very substantially when the diacritic was not presented, indicating an excessive reliance on the diacritic. Although this outcome is justifiable on the basis of the reliability and validity of the stress diacritic, it stands in contrast to behavioral data showing that Greek readers are not affected by the lack of a stress diacritic (Protopapas, Gerakaki, & Alexandri, 2007; Protopapas & Gerakaki, in press). Pre-training the phonological layer might produce an improved fit to human performance by reinforcing stress vowel con-

nections in word representations. A connection from the orthographic directly to the phonological layer (Zorzi et al., 1998) might also improve performance on unstressed words.

In this preliminary investigation, two sublexical properties were found to affect reading times: word length (measured in letters or, for nonwords, in syllables) and bigram probability. In a review of factors affecting visual word recognition, Balota et al. (2006) reported significant effects of word length for low frequency words and nonwords. The effect on both words and nonwords in our results may be due to the relatively few repetitions of each word during training. This renders words effectively low-frequency, because they did not have many opportunities to affect the connection weights. An alternative or complementary explanation may relate to a fine grain of graphophonemic representation, which is expected for a language with high feedforward consistency. Reliance on a fine grain can lead to stronger length effects as more graphemic units must be individually mapped. As a reviewer pointed out, this might also explain the lack of frequency and lexicality effects. It remains to be investigated whether evidence for larger units of graphophonemic mapping may accumulate with higher ratios of trials to word types in the training procedure.

Nevertheless, the significant word length effect seems to run counter to common expectations regarding connectionist models. According to Rastle & Coltheart (2006), word length effects should not be exhibited by single-route connectionist models, because entire words are read in parallel and not serially, grapheme-by-grapheme, as in some dual-route models. Our results, although preliminary, are inconsistent with this prediction, showing that word length effects are possible in parallel distributed processing models, even for the highly consistent mappings of the Greek orthography. This finding may depend on a large range of word lengths, as imposed by the multisyllabic input and by the stimulation of more phonological attractors when more letters appear at the input, in part due to dispersion. Thus, our model sheds light on a long-standing issue in modeling reading aloud, which was not possible to address with previous models dealing only with monosyllabic words.

Word frequency is one of the most important variables affecting word reading performance in English (Balota et al., 2006). Our model was not affected by word or syllable frequency in this preliminary investigation, which may be attributed to the low token-to-type ratio that renders trained words effectively low frequency. Balota et al. noted that low frequency words exhibit larger effects of sublexical regularity, such as bigram frequency, compared to high frequency words. This was borne out in our model and may be related to the dispersion necessitated by our input representation. Specifically, as the same letters appear at different positions, the model is exposed to input bigrams more consistently than to words with larger common parts.

On the other hand, the absence of expected transparency (mapping consistency) effects warrants further investigation. Orthographic neighbors were also expected to affect reading performance, however the situation with neighbors

may differ substantially from English because most Greek words have few or no neighbors, (mean neighborhood size was 1.69), perhaps due to their overall greater length.

In conclusion, this paper presents a computational model of reading aloud that can read Greek multisyllabic words and nonwords, using a novel orthographic input representation that includes stress marking. Critically, orthographic input was not pre-syllabified, whereas phonological output was. Preliminary tests indicate that the model reads words and nonwords with reasonably high accuracy, assigns stress correctly based on diacritic information, and produces effects of word length, previously thought incompatible with parallel processing, but no effects of frequency, which are large and robust in human data and other models. Further tests and elaboration will take place as comparable human data for Greek become available.

Acknowledgments

We thank Jason D. Zevin for the source code used to implement the model and Efthymia C. Kapnoula for the word and nonword testing stimuli.

References

- Balota, D. A., Yap, M. J., & Cortese, M. J. (2006). Visual word recognition: The journey from features to meaning (A travel update). In M. Traxler & M. A. Gernsbacher (Eds.) *Handbook of psycholinguistics (2nd ed.)*.
- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review, 108*, 204–256.
- Chalamandaris, A. Raptis, S., & Tsiakoulis, P. (2005). Rule-based grapheme-to-phoneme method for the Greek. In *INTERSPEECH-2005*
- Harm, M. W., & Seidenberg, M. S. (1999). Phonology, reading acquisition and dyslexia: Insights from connectionist models. *Psychological Review, 106*, 491–528.
- Harm, M. W., & Seidenberg, M. S. (2004). Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review, 111*, 662–720.
- Monaghan, P., Arciuli, J., & Seva, N. (2008). Constraints for computational models of reading: Evidence from learning lexical stress. In *Proceedings of the 30th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum.
- Pagliuca, G., & Monaghan, P. (in press). Discovering large grain-sizes in a transparent orthography: insights from a connectionist model of Italian word naming. *European Journal of Cognitive Psychology*.
- Pagliuca, G. & Monaghan, P. (2009). A connectionist model of reading for Italian. In J. Mayor, N. Ruh, & K. Plunkett (Eds.), *Connectionist models of behaviour and cognition II* (pp.301–312). Singapore: World Scientific.
- Perry, C., Ziegler, J. C., & Zorzi, M. (2007). Nested incremental modeling in the development of computational theories: The CDP+ model of reading aloud. *Psychological Review, 114*, 273–315.
- Petrounias, E. V. (2002). *Neoelliniki grammatiki kai sigkritiki analisi, tomos A: Fonitiki kai eisagogi sti fonologia [Modern Greek grammar and comparative analysis, Vol A: Phonetics and introduction to phonology]*. Thessaloniki: Ziti.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review, 103*, 56–115.
- Protopapas, A. (2006). On the use and usefulness of stress diacritics in reading Greek. *Reading & Writing: An Interdisciplinary Journal, 19*, 171–198.
- Protopapas, A., Gerakaki, A., & Alexandri, S. (2007). Sources of information for stress assignment in reading Greek. *Applied Psycholinguistics, 28*, 695–720.
- Protopapas, A., & Gerakaki, S. (in press). Development of processing stress diacritics in reading Greek. *Scientific Studies of Reading*.
- Protopapas, A., & Nomikou E. (2009). Rules vs. lexical statistics in Greek nonword reading. To appear in the *Proceedings of the 31st Annual Conference of the Cognitive Science Society*.
- Protopapas, A., & Vlahou, E. L. (in press). A comparative quantitative analysis of Greek orthographic transparency. *Behavior Research Methods*.
- Rastle K, Coltheart M (2006), "Is there serial processing in the reading system; and are there local representations?", In S. Andrews (Ed.), *From inkmarks to ideas: Current issues in lexical processing*. Hove: Psychology Press.
- Rastle, K., & Coltheart, M. (2000). Lexical and nonlexical print-to-sound translation of disyllabic words and non-words. *Journal of Memory & Language, 42*, 342–364.
- Revithiadou, A. (1999). *Headmost accent wins: Head dominance and ideal prosodic form in lexical accent systems* [LOT Dissertation Series 15 (HIL/Leiden University)]. The Hague: Holland Academic Graphics.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review, 96*, 523–568.
- Seidenberg, M., & Plaut, D. (2006). Progress in understanding word reading: Data fitting versus theory building. In S. Andrews (Ed.), *From inkmarks to ideas: Current issues in lexical processing*. Hove: Psychology Press.
- Spencer, K. A. (2007) Predicting children's word-spelling difficulty for common English words from measures of orthographic transparency, phonemic and graphemic length and word frequency. *British Journal of Psychology, 98*, 305–338.
- Zevin, J. D., & Seidenberg, M. S. (2006). Simulating consistency effects and individual differences in nonword naming: A comparison of current models. *Journal of Memory and Language, 54*, 145–160.
- Zorzi, M., Houghton, G., & Butterworth, B. (1998). Two routes or one in reading aloud? A connectionist dual-process model. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 1131–1161.