The effect of phonological encoding on word duration: Selection takes time

Duane G. Watson[1], Andrés Buxó-Lugo[1], & Dominique C. Simmons[2]

[1]University of Illinois Urbana-Champaign

[2]University of California Riverside

Manuscript Dated: 11/4/13

Word Count: 3742 (excluding references)

Mailing address:

603 E. Daniel St.

Department of Psychology

University of Illinois Urbana-Champaign

Champaign, IL 61820

E-mail: dgwatson@illinois.edu

Running head: PHONOLOGICAL ENCODING AND DURATION

Abstract

In this chapter, we investigate whether the process of phonological encoding plays a role in determining the duration of a word. We explore whether points of complexity in word production as predicted by a simple recurrent network also predict points within a word at which speakers slow down. Simple recurrent networks were trained to produce two different words under two conditions: in the first condition the two words in the sequence overlapped in their initial morphemes (e.g. layover layout) and in the second condition the words overlapped in their final morpheme (e.g. overlay outlay). The network experienced the most error for words that overlapped initially and at points of word non-overlap. Participants who produced these same sequences in a repetition task exhibited lengthening at points of complexity predicted by the network. We propose that lengthening may be partly a result of the phonological encoding system needing processing time.

It is a well-known phenomenon that speakers lengthen words that are new, informative, or not predictable in a conversation and shorten words that are given, predictable or non-informative (e.g. Aylett & Turk, 2004; Bell et al., 2009; Fowler & Housum, 1987; Jurafsky, 2001; Lam & Watson, 2010; Pluymaekers et al., 2005 and many others). A puzzle for linguists, psychologists, and computer scientists who are interested in prosody is understanding why.

There are two varieties of explanations. One is that speakers lengthen and shorten words to facilitate robust communication with listeners. This idea has been described within formal frameworks like the Uniform Information Density Hypothesis (e.g. Jaeger, 2010), and the Smooth Signal Hypothesis (Aylett & Turk, 2004): speakers lengthen linguistic information with high information content and shorten words with low information content to create a uniform information density across utterances. The other explanation is that the duration of words partly reflects the complexity of underlying production processes. Speakers produce words that are new or informative with longer duration because those words are actually more difficult to say. The extra time provided by lengthening the segments facilitates the production process.

It is important to note that these two explanations are not incompatible. It is possible that duration choices facilitate the workings of mechanisms that are engaged in production while at the same time optimizing word length for robust communication. However, a challenge for both of these approaches is mapping out the underlying mechanisms.

The current chapter explores the algorithms that underlie production-centered theories of reduction and lengthening. There have been some proposals for how reduction and lengthening might facilitate production (e.g. see Bell et al. 2009, Kahn & Arnold, in press, Kahn & Arnold, 2012), though typically the mechanism is framed in terms of activation of the routines associated with production. If a word has been recently produced or is highly predictable, its resting activation will be higher, and consequently, it will be easier to produce, and the word will be shortened. Similarly, because new words will have lower activation, the increased effort required for articulation results in longer production times. A drawback of this type of explanation is that it remains unspecified as to why lengthening a difficult word (or reducing a highly activated word) would facilitate language production. If lengthening is linked to planning difficulty, why does it not occur before the critical target word? Once one begins to utter a word, presumably its meaning and lemma have already been accessed. What benefit could a speaker derive from lengthening a new word once articulation has already begun?

The answer may lie in theories of phonological encoding. In some models of word production, phonological selection is a serial process (Sevald & Dell, 1994; although see O'Seaghda & Marin, 2000). Once a word is accessed, phonemes are accessed in an order that corresponds with the order in which they appear in a word, starting with phonemes at the beginning of the lexical item. There is empirical support for this type of architecture. Sevald & Dell (1994) found that rapid repetition of two alternating words was faster when those words shared their rhymes ("TICK" vs "PICK") than when they shared onsets ("PICK" vs "PIN") (see Jaeger et al., in press, for similar effects of phonological overlap on lexical selection). These results can be accounted for within an interactive activation model like the Dell (1986) model in which low-level phonemic representations send feedback to higher-level lexical representations. In such a model, the shared onset activates both words, which increases competition between the lexical items and inhibits the correct selection of the target. In contrast, words that share rhymes are not burdened by inhibition early in the selection process, which facilitates production.

If phonological encoding is a serial process, or at least a process that is not entirely completed at the point of articulation, this may explain why words that are new are lengthened. Lengthening could provide more time for phonological selection to take place at the point of articulation. Similarly, word reduction could be the result of faster phonological selection. There are empirical data that suggests these effects are in fact driven by phonological processes. Kahn & Arnold (2012) found that both non-mentioned, conceptually given words and mentioned words are reduced, but mentioned words exhibit greater reduction. Similarly, Lam & Watson (in press) found that repeated words, but not repeated referents, lead to reduction. Although the results in Kahn & Arnold (2012) and Lam & Watson (in press) do not by themselves suggest that a serial production processes underlies these effects, they do suggest that these effects originate at the phonological or articulatory level.

In this paper we explore whether the dynamics of a phonological production system that serially encodes linguistic information can explain changes in duration in word production. The strategy we use is to first understand whether a serial selection model predicts complexity at varying points within a word and across words. Then we test to see whether English speakers' durational choices match predicted points of complexity by the model. If predicted points of complexity and lengthening overlap, it will suggest that duration effects might be linked to to phonological encoding processes.

We use a model inspired by Dell, Juliano, & Govindjee (1993). It is a simple recurrent network, originally designed to model speech errors. This model was used for two reasons. The first is that it allowed us to easily encode phonological selection as a serial process that occurs across time. As in all SRNs, Dell et al.'s model has a set of context units that encodes activation of hidden nodes on previous time steps. The second motivation for using this model was that it allowed us to test whether representational similarity across words impacts difficulty of production while making minimal assumptions about the architecture of the model.

As in Sevald & Dell (1994), the model was trained to produce words in which the output overlaps in its initial part (e.g. "layover layout") or overlaps in its final part ("overlay outlay"). We used words with morphological overlap instead of overlap in sub-syllabic components like the rime and onset (as in Sevald & Dell, 1994). This was done to increase the amount of overlap across words in order to amplify the size of any potential effect that this might have on word production in our human production data. Manipulating morphological overlap also has the added advantage of simplifying the learning goals of the model: rather than learning mappings between a lemma and a long string of phonological features or phonemes, the model learns a simple mapping between a lemma and two parts of a word, allowing us to focus our question on how linguistic overlap generally impacts production. Finally, by making the unit of overlap a morpheme, we can more easily measure duration differences across words in human productions.

The prediction from Sevald & Dell's (1994) results is that words that overlap initially should be more difficult to produce than those that overlap finally. Critically, we will see where *within* the words the model predicts the greatest point of complexity, and determine whether these predictions correspond with speaker duration preferences.

Models

The architecture of the model was similar to that used by Dell et al. (1993). The primary components are illustrated in Figure 1.

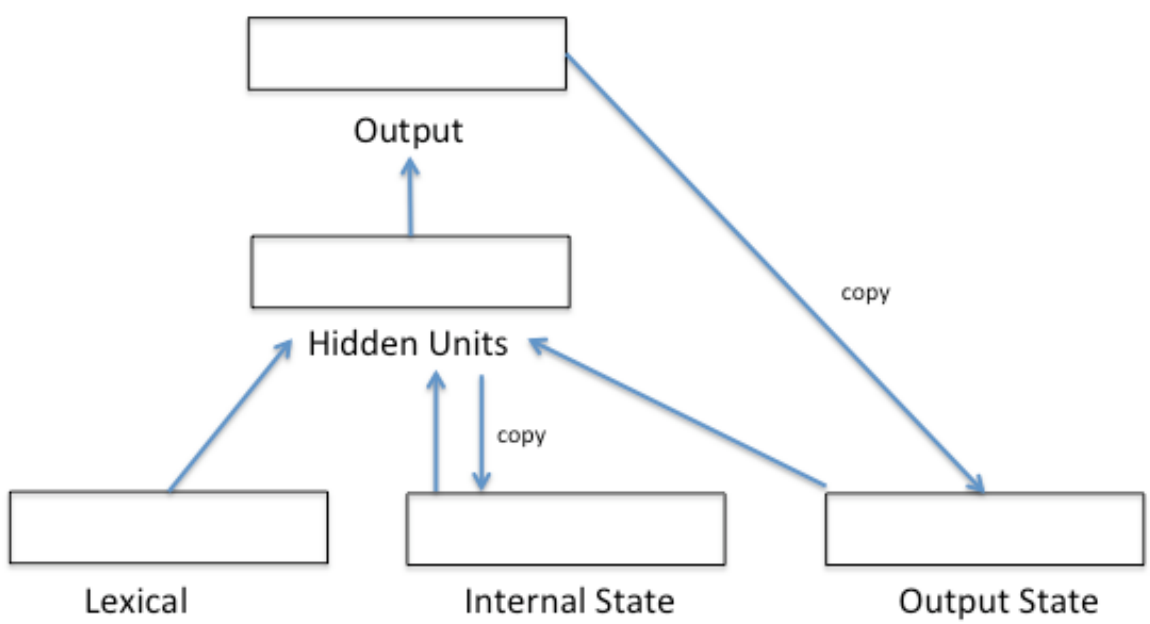

Figure 1. The architecture of the simple recurrent network.

The network consists of an input layer that represents the lexical context, a hidden layer, an output layer that generates a morpheme based on the lexical input it is receiving, and two context layers that represent the state of the hidden layer on the previous time step and the state of the output layer on the previous time step.

On each cycle of the model, activation propagates from the lexical layer to the hidden layer and from the hidden layer to the output layer. Activation to each node was computed using the logistic activation function. After each cycle, the activations of nodes in the hidden layer are copied to the internal state layer. The activations of nodes in the output layer are copied to the output state layer. On the next cycle, activation from both the internal state layer and the output state layer propagate their activation along with that of the lexical layer to the hidden units. This allows for the state of the units in the hidden and output layers on previous cycles to influence processing of hidden units on the current cycle, giving the model a memory (see Elman, 1990; Jordan 1986). Dell et al. (1993) tested versions of the model in Figure 1 with both internal state and output state layers and with just an internal state layer. They found that errors produced by the former more closely matched the errors produced by speakers, so we use the same architecture here.

The model was trained using the back-propagation learning algorithm (Rumelhart, Hinton, & Williams, 1986). The input layer consisted of 2 nodes, one for each lexical item to be produced (e.g. layout vs. layover). The hidden layer consisted of 7 nodes, as did the internal state layer. The output layer consisted of four nodes, as did the output state layer. The four nodes of the output layer corresponded to each of the morphemes in the two word vocabulary (e.g. lay, over, out) as well as a node that corresponded with a word boundary.

The training vocabulary consisted of two words. The model was trained to produce the two words in alternation (layover-layout-layover-layout-etc…). Models were trained on two word vocabularies in order to determine, in general, how final and initial overlap impact production difficulty. Although we would expect to find similar effects in models with larger vocabularies, by examining models with two words, we can focus specifically on effects of overlapping representations on production rather than effects of other factors such as interactions across lexical items or model memory constraints. In addition, a two word vocabulary allowed the learning phase of the model to more closely match the task performed by participants, which we discuss below, a production task in which two lexical items are produced in sequence.

On each cycle, the input node corresponding to the target word was activated. This activation occurred across three time steps to produce the two morphemes of the word and the word boundary (e.g. lay, over, word boundary). Training ended after 200 epochs, which included 2 productions of the 2 compound words each (i.e. "layout layover layout layover" 200 times).

Two types of models were trained. One group of models was trained to produce two words that overlapped in their initial morphemes (e.g. layout and layover). Another group of models was trained to produce two words that overlapped in their final morphemes (e.g. outlay and overlay). At test, the models were given a target two-word sequence to produce. We used the mean summed square error of the output nodes as an indicator of overall model difficulty in producing each of the morphemes. Figure 2 displays the average summed squared error for ten models trained on words that overlap initially and for ten models trained on words that overlap finally.

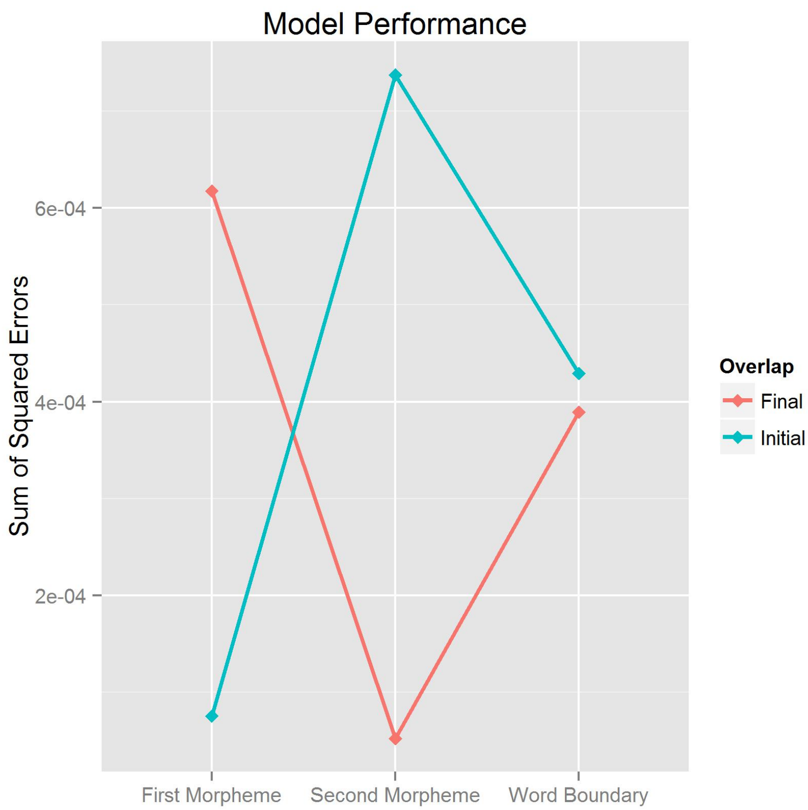

Figure 2. The average summed squared error across output nodes for ten models trained on words that overlap in their initial morpheme (layout-layover) and ten models trained on words that overlap in their final morpheme (outlay-overlay).

The overall pattern replicates what one would expect from Sevald & Dell's (1994) results: words that overlap in their initial segments are more difficult to produce than those that overlap in their final segments: summing the squared error over the three regions yields in error of 1.1879 for the initial overlap condition and 1.0061 for the final overlap condition. The second thing to note is that both models predict more difficulty at points at which the words do not overlap than at points at which they do, predicting that the most distinctive part of the word should be the one that is the most difficult to generate for speakers.

As in the model proposed by Sevald & Dell (1994), the serial nature of phonological encoding readily explains this pattern. Representational similarity creates more difficulty when it occurs earlier in the word. In the SRN, retrieved material on the previous cycle serves as a partial cue for retrieving material at the present cycle. Thus, the input to the hidden layer for "layover" and "layout" is similar at the points of the second morpheme, and this representational similarity leads to more difficulty in producing it. In contrast, for words that overlap finally, this representational similarity does not occur until the word is near completion, and thus, does not create interference. Thus, the difference in difficulty in producing words that overlap initially and finally is the result of competing representations between words as suggested by Sevald & Dell (1993) in their interactive model.

Thus, a production model based on a simple recurrent network architecture mirrors the performance of human speakers (see Sevald & Dell, 1994). However, the present question is whether the relative duration of points within a word reflects this speaker difficulty. If it is the case that the sequential nature of phonological encoding influences the timing of word production, we would expect points at which the model has difficulty producing a word like "layout" to predict the relative duration of points within the word. For both sets of models, producing the non-overlapping part of the word resulted in the most error. This is likely, in part, a frequency effect: the overlapping morpheme was far more frequent in the input than the non-overlapping morpheme. In addition, non-overlapping morphemes were the most distinctive parts of the words, so the contexts that preceded them (and their inputs to the hidden layer) tended to be similar resulting in increased difficulty,, another interference effect.

Thus, two effects are predicted: 1) Speakers should produce words that overlap initially with longer duration than words that overlap finally and 2) We expect greater slow downs on parts of the word that do not overlap in the two conditions. In the experiment described below, speakers produced alternating word sequences that overlapped either initially or finally such as "layout layover layout layover…" or "outlay overlay outlay overlay". We investigated how this overlap affected word duration across these two conditions.

Method

Participants

Fifteen undergraduates from the University of Illinois Urbana-Champaign participated in the study for class credit.

Materials

Because we were interested in the effects of word overlap on duration, it was critical to control for other factors known to affect word duration such as lexical stress, metrical stress, and phonological context. To control for the phonological context of the target words, we created word sets in which reversing the order of the morphemes in the compound produced English words that overlapped either initially or finally such as in (1) below, so that the strings produced across conditions were matched and varied only in their order.

1a) layover layout

1b) overlay outlay

There were a total of six items within each condition. The words differed across conditions, but as in (1), these words were matched with respect to the morphological components that were used.

In addition, unlike in the modeling data, we could not simply compare performance across the two conditions. Word duration is affected by metrical stress, syllable position, and other potential confounding factors. Thus, target words, such as "layout" and "layover", were compared to baseline conditions that included one of the critical words. This word was paired with a compound with which it did not overlap morphologically. Thus, for the string "layout layover", the duration of "layout" in critical trials was compared to "layout" in a baseline condition "layout handover". Similarly, a baseline was constructed for the other member of the target set: the baseline for "layover" was "layover handout". Baseline conditions were also constructed for the final overlap conditions. Thus, the data presented below represents the difference in duration between the target words in initial and final overlap conditions and their respective baseline conditions. All the conditions and their baselines are listed in (2) in an example item.

2a) Initial overlap: layover layout

2b) Final overlap: overlay outlay

2c) Initial baseline 1: **layover** handout

2d) Initial baseline 2: **layout** handover

2e) Final baseline 1: **overlay** handout

2f) Final baseline 2: **outlay** handover

A within-participant design was used. In addition, all participants produced both conditions for each item as well as their associated baseline conditions. This was done to reduce potential inter-speaker differences in pronunciation and speech rate. These stimuli were presented to subjects with fifty distractor items that consisted of unrelated compounds (e.g. horseshoe nightshade, hindsight staircase).

Two factors were counterbalanced across participants. One was the order in which the critical items were presented to participants (layover layout vs. layout layover), which yielded two lists. Critical and baseline items were randomized within these two lists. In order to counterbalance order of presentation, two more lists were constructed with items presented in reverse order, yielding a total of four lists.

Procedure

Each trial consisted of a single word pair that participants were told to say aloud as quickly as possible and as many times as possible without errors. Participants were given eight seconds to speak. Participants completed several practice trials before proceeding to the main body of the experiment, which consisted of a total of 86 trials.

To analyze durations, each morpheme of each compound was labeled in Praat (Boersma & Weenink, 2012), a speech-analysis platform. Morpheme duration was automatically extracted using a script.

Results

Out of 15,316 morphemes, a total of 323 word errors were made. These errors were excluded from the analysis, leaving 97.89% of the data in the analysis. Errors included disfluencies within a word, coughs during production, producing an incorrect word, producing only part of the target compound, and producing a correct word but with neighboring incorrect words. The remaining data were analyzed using multilevel linear mixed effects models with fixed effects of trial type (Target vs. Baseline), location of the morpheme in the word (first vs. second), and where in the word the morphological overlap occurred (final vs. initial morpheme). All three factors were centered. Reported p-values were obtained by assuming that, given the number of observations, the t-distribution approximated a z-distribution. Following the recommendations of Barr et al. (2012), the maximal random effects structure was used.

Table 1.

Fixed Effect Estimates for Multi-Level Model of Participant Durations.

	Estimate	Standard Error	t value

Intercept	0.0.2596	0.0079	32.80
Overlap Location	0.0022	0.0096	0.23
Target (vs. Baseline)	0.0021	0.0047	0.46
Morpheme Location	0.0472	0.0107	4.42
Overlap * Target	0.0095	0.0099	0.96
Overlap * Morpheme Location	0.0155	0.0194	0.80
Target * Morpheme Location	0.0004	0.0118	0.03
Overlap * Target * Morpheme Location	0.0471	0.0239	1.97

The fixed effects are presented in Table 1. Overall, there was a bias towards producing the second morpheme with longer duration than the first. This was true in all conditions, and probably reflects a metrical structure imposed on the words by the participants given the repetitive nature of the task.

Critically, there was a reliable three-way interaction between trial type, morpheme position, and overlap ($t=1.97$, $p < 0.05$). The morpheme durations are displayed in Figure 3.

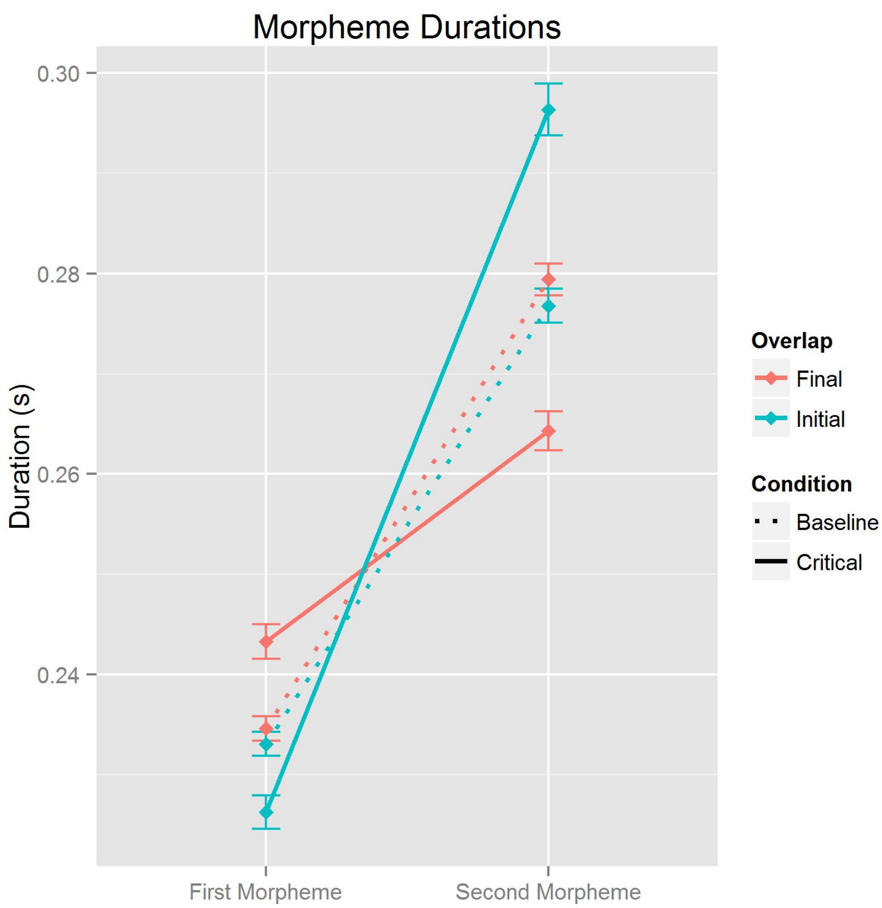

Figure 3. The mean duration in seconds of morphemes across trial. Error bars show standard errors.

In the condition in which morphemes overlap initially, target durations were longer than the corresponding baseline condition at the second morpheme. In contrast, in the condition in which morphemes overlap finally, target durations were shorter than the corresponding baseline condition at the second morpheme. The opposite pattern is true at the first morpheme: the initial overlap condition is shorter than its corresponding baseline control while the final overlap conditions is longer than its corresponding baseline control. To highlight these differences, Figure 4 displays a graph of the durations of initial and final overlap conditions with the duration of their baselines subtracted.

Figure 4. A difference score between the durations of the morphemes in the initial and final overlap conditions and their corresponding baseline conditions.

Note that the relative durations of the first and second morpheme in both conditions closely match the error predictions of the model in Figure 2: the model experiences the most difficulty on the portion of the word that does not overlap, and we find that this matches the human duration data.

Finally, contrary to predictions, there was no overall difficulty effect of overlap. It was not the case that the overall difference between the initial overlap condition and its baseline (M= 0.0036s) was significantly different than the final overlap condition and its baseline (M= -0.0051s) ($t=.96$, $p>.005$). The effect was numerically in the right direction though it was not reliable. Previous results have found that initial overlap leads to shorter durations than final overlap (e.g. Sevald & Dell, 1994, O'Seaghda & Marin, 2000). It is possible that this effect did not reach significance in the current study because of power: only six items could be used in each condition because there are very few compound pairs in English that overlap initially and, when reversed, overlap finally and still yield real English words. These design constraints may have limited our ability to detect an effect if it was present.

Discussion

This chapter began with a puzzle: if variability in the duration of a word is linked to production difficulty, and lengthening difficult words facilitates production, why does this lengthening occur during word production rather than before it?

The goal of this chapter was to demonstrate that production processes involved in phonological selection can provide a partial explanation for duration differences both within and across words. If phonological selection is a serial process that inevitably varies in complexity at different time points, lengthening at points of uncertainty might facilitate production by giving the system more time to converge on selecting the correct phonemes. The simple recurrent network presented above predicted that words that overlap initially should be more difficult to produce than words that overlap finally. It also predicted that production should be more difficult in the regions of the word that do not overlap. We found that participants that produced word pairs in contexts similar to the model slowed down in exactly the points at which the model predicted production difficulty. The fact that points of complexity correspond with lengthening suggests that some durational choices by speakers may be attributable to the process of phonemic encoding.

As discussed above, we did not replicate the effect of overlap found in previous study (e.g. Sevald & Dell, 1994; O'Seaghda & Marin, 2000), though this was predicted by the model. This may have been due to insufficient power. However, it is encouraging that the SRN correctly accounts for the findings of previous work: initial overlap leads to more difficulty than final overlap. Furthermore, the model correctly predicts that the non-overlapping morphemes of the compounds should be produced with longer durations than overlapping morphemes.

Note that we are not arguing that production constraints are the only factors that affect word duration. Factors like word or lemma frequency, speech rate, and communicative factors such as those outlined in Aylett & Turk's (2004) Smooth Signal Hypothesis almost certainly contribute to the duration of a word. Nevertheless, complexity in the production system could help explain why at least some of these factors contribute to changes in word durations.

Another question is understanding the level of production at which these duration effects arise. Above, we attribute the duration effects to mechanisms linked to serially ordering phonological information, however, these data are also consistent with complexity in the ordering of any sub-lexical linguistic production process (e.g.

phonological, articulatory, morphological, or syllabic representations). Furthermore, these data are also compatible with certain types of non-serial production processes. That is to say, these data are compatible with sub-lexical linguistic production routines not reaching completion or occurring at different stages. The process of selection might only be partially serial, and still yield the types of differential lengthening we see across words. For example, the phonemes that are most highly activated might be selected first while phonemes that are less activated are only selected at a later stage. Both processes might occur during articulation, but the system itself is not entirely serial. Such an architecture is consistent with the larger point being made here: that duration choices allow time for linguistic selection, but they do not necessarily assume a fully serial architecture. We leave the question of what level of representation and the degree of seriality in the production system open to future investigation.

Finally, it is important to note that the SRN presented above is not meant to be a model of the representations that are engaged in language production. Although the model demonstrates that the difficulty of ordering linguistic information is sensitive to overlap between compounds and that these map onto speakers' durational choices, it does not necessarily model the actual mechanisms that are engaged in language production. One approach for developing a process model would be to adapt the Dell (1986) model so that it captures some of serial effects in this chapter, as well as the effects reported in Sevald & Dell (1996) and O'Seaghda & Marin (2000). This might serve as a useful next step in understanding how production algorithms lead to differences in duration.

Overall, the SRN and the behavioral data point towards a link between production processes and duration. Although there are claims in the literature that such a link exists, up until now, there has been relatively little work specifying exactly how lengthening and reduction are linked to the process of speaking. The work here represents a first step in spelling out the mechanisms that underlie this link.

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.

Aylett, M. & Turk, A. (2004). The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language and Speech*, 47(1), 31-56.

Barr, D.J., Levy, R., Scheepers, C., & Tily, H.J. (in press). Random effects structure in mixed-effects models: Keep it maximal. *Journal of Memory and Language*.

Bell, A., Brenier, J.M., Gregory, M., Girand, C., & Jurafsky, D. (2009). Predictability effects on durations of content and function words in conversational English. *Journal of Memory and Language*, 60, 92-111.

Boersma, P. & Weenink, D. (2012). Praat: doing phonetics by computer [Computer program]. Version 5.3.32, retrieved 17 October 2012 from http://www.praat.org/

Bybee, J. L. (2001). *Phonology and language use,* Cambridge University Press: Cambridge.

Dell, G.S. (1986). A spreading activation theory of retrieval in sentence production. *Psychological Review*, 93, 283-321.

Dell, G.S., Juliano, C., & Govindjee, A. (1993). Structure and content in language production: A theory of frame constraints in phonological speech errors. *Cognitive Science*, 17, 149-195.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.

Jaeger, T.F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology,* 61, 23-62.

Jaeger, T.F., Furth, K., & Hillard, C. (in press). Phonological overlap affects lexical selection during sentence production. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

Jordan, M.I. (1986). Serial order: A parallel distributed processing approach. Tech. Rep. No. 8604. San Diego: University of California, Institute for Cognitive Science.

Jurafsky, D., Bell, A., Gregory, M., and Raymond, W.D. (2001). Probabilistic Relations between Words: Evidence from Reduction in Lexical Production. In J. Bybee and P.

Hopper (eds.). *Frequency and the emergence of linguistic structure.* Amsterdam: John Benjamins. 229-254

Fowler, C.A. & Housum, J. (1987). Talkers' signaling of "new" and "old" words in speech and listener's perception and use of the distinction. *Journal of Memory and Language*, 26, 489-504.

Kahn, J. & Arnold, J.E. (in press). Speaker-internal processes drive durational reduction. *Language and Cognitive Processes*.

Kahn, J. & Arnold, J.E. (2012). A processing-centered look at the contribution of givenness to durational reduction. *Journal of Memory and Language, 67*, 311-325.

Lam & Watson (in press). Repetition Reduction? Journal of Experimental Psychology: Learning, Memory, & Cognition.

Lam, T. Q. & Watson, D. G. (2010). Repetition is easy: Why repeated referents have reduced prominence. *Memory & Cognition*, *38*, 1137-1146.

Marslen-Wilson, W. (1987). Functional parallelism in spoken word recognition. *Cognition*, 25, 71-102.

McClelland, J.L. & Elman, J.L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18, 1-86.

O'Seaghda, P.G. & Marin, J.W. (2000). Phonological competition and cooperation in form-related priming: Sequential and nonsequential processes in word production. *Journal of Experimental Psychology: Human, Perception, and Performance*, 26, 57-73.

Pluymaekers, M., Ernestus, M., & Baayen, R.H. (2005). Lexical frequency and acoustic reduction in spoken Dutch. *Journal of the Acoustical Society of America*, 118, 2561-2569.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propogation. In. D.E. Rumelhart, & J.L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol 1., pp. 318-362). Cambridge, MA: MIT Press

Sevald, C.A. & Dell, G.S. (1994). The sequential cuing effect in speech production. *Cognition*, 53, 91-127.

Author's Note

We would like to thanks John Hummel and Gary Dell for their comments and input. This work was supported by NIH grant R01 DC008774 and a grant from the James S. McDonnell foundation.