Running head: FEEDBACK IN SPEECH PRODUCTION

Articulatory Evidence for Feedback and Competition in Speech Production

Corey T. McMillan and Martin Corley*

University of Edinburgh

School of Philosophy, Psychology, & Language Sciences

Robin J. Lickley

Queen Margaret University

Speech & Hearing Sciences

*Address correspondence to:

Martin Corley

Psychology

School of Philososophy, Psychology, and Language Sciences

University of Edinburgh

7 George Square

Edinburgh EH8 9JZ, UK

(tel) +44 131 650 6682; (fax) +44 131 650 3461; Martin.Corley@ed.ac.uk

Abstract

Speech error investigations have traditionally relied on the perceptual categorization of responses, limiting analyses to small subsets of the data identified as appropriate 'errors'. Here we report an experimental investigation of slips of the tongue using a Word Order Competition (WOC) paradigm in which context (entirely nonlexical, mixed) and competitor (whether a possible phoneme substitution would result in a word or not) were crossed. Our primary analysis uses electropalatographic (EPG) records to measure articulatory variation, and reveals that the articulation of onset phonemes is affected by two factors. First, onsets with real word competitors are articulated more similarly to the competitor onset than when the competitor would result in a nonword. Second, onsets produced in a nonlexical context vary more from the intended onset than when the context contains real words. We propose an account for these findings that incorporates feedback between phonological and lexical representations in a cascading model of speech production, and argue that measuring articulatory variation can improve our understanding of the cognitive processes involved in speech production.

Articulatory Evidence for Feedback and Competition in

Speech Production

Anyone who tries to say a phrase such as *Bob flew by Bligh Bay* will need little convincing that in speech, there are often substantial differences between intention and articulation. These differences are not arbitrary, however: When individual words are mispronounced, the resulting speech is more likely to contain real words than would be created by chance insertions, deletions, or substitutions of phonemes (Dell & Reich, 1981; Nooteboom, 2005; but see del Viso, Igoa, & Garcia-Albea, 1991; Garrett, 1976). The bias towards real words, or *lexical bias*, reveals that slips of the tongue cannot be attributed solely to inaccurate motor programming of the tongue, but must reflect planning aspects of the speech production process.

Two different loci of the lexical bias effect have been proposed within current models of speech production: The effect has been attributed either to the self-monitor, or to feedback between levels of the speech production system. Many proponents of the self-monitoring account favor a feedforward view of speech production, in which processing at each level, with very few exceptions, leads to the selection of only a single unit which is passed on to the next level (e.g., Levelt, 1989). The self-monitor has access to speech plans or "inner speech" (Levelt, 1983), and can employ a lexicality criterion such that ill-formed speech plans which would have resulted in nonwords are edited out prior to articulation (Baars, Motley, & MacKay, 1975; Levelt, 1989; Motley, Camden, & Baars, 1982). Proponents of the feedback account suggest instead that lexical bias is the consequence of feedback from phonemic to lexical representations. Since nonwords are not represented in the lexicon, misactivated phonemes can only increase the activations of real word representations. Activation of real word representations in turn boost the activation of corresponding phonological representations. This boost, which occurs for real words,

but not nonwords, makes it more likely that a real word will be produced in error (e.g., Dell, 1986).

A primary source of evidence used to distinguish the two accounts is the SLIP task, a speech error elicitation paradigm originally developed by Baars and colleagues (Baars et al., 1975; Baars & Motley, 1976). Baars et al. (1975, Experiment 2) embedded target word pairs into context lists which consisted either entirely of nonwords, or of a mixture of words and nonwords. In each trial, participants silently read several sequentially presented biasing items (e.g., *keet fime*). On occasion they were asked to repeat aloud a target with reversed onsets (e.g., *feep kive*). Targets were designed so that exchange errors would result in real words (e.g., *feep kive* → "keep five"; henceforth we represent the intended utterance with *italics*, and the spoken utterance in quotation marks) or nonwords (e.g., *feeb kise* → "keeb fise"). In the mixed context, Baars et al. observed more phoneme exchanges that resulted in real words than nonwords, establishing experimental evidence for the lexical bias. In the nonword context, however, the exchange levels did not differ for real word and nonword outcomes. Baars et al. reasoned that in conditions where real words were never encountered (nonlexical context), there was no need for the self-monitor to make use of a lexicality criterion to monitor and edit the speech plan. Since there is no clear reason that feedback should be modulated by a context manipulation, Baars et al.'s finding has often been taken as strong evidence for a monitoring account (e.g., Levelt, 1989; Levelt, Roelofs, & Meyer, 1999; Roelofs, 2004).

However, more recent evidence from the SLIP task has been more equivocal. For example, Hartsuiker, Corley, and Martensen (2005) also observed a lexical bias in the mixed context but a different pattern in the nonlexical context and Humphreys (2002) observed a lexical bias in both contexts. These findings have been variously interpreted as support for a hybrid model in which both feedback and self-monitoring contribute to the pattern of results, or as support for a feedback-only model. What is most striking,

however, is that the effect of context on slips of the tongue remains far from clear across studies. In fact, the reported numbers of exchanges are very low: In total, 8.2% of responses were described as full exchanges by Baars et al. (1975); Humphreys (2002) reports 2.7%; the figures for Hartsuiker et al. (2005) are 4.1% (Experiment 1) and 2.3% (Experiment 2). With low numbers of occurrences and empty cells, the data may be susceptible to noise (see also Nooteboom & Quené, in press).

One reason for the low numbers of reported exchanges may be that participants' responses are categorized, by deciding whether each response is an appropriate 'phoneme exchange error' or not. As a consequence, a number of mispronunciations recorded during the course of each experiment are discounted from further analysis. In many cases, the responses that are categorized as 'errors' but are not deemed appropriate for analysis greatly outnumber the reported exchanges (for example, Hartsuiker et al. (2005) report 10.4% and 13.9% 'other' errors for Experiments 1 and 2 respectively). Moreover, the boundaries between categories such as 'exchange', 'other error', and 'correct' may be difficult to determine: Frisch and Wright (2002) highlighted that, when transcribing speech, listeners perceive sounds that are intermediate between categories as belonging to one category (Liberman, 1997) and that distorted speech sounds may be subject to phoneme restoration effects (see Samuel, 1996, for a review).

In fact, the existence of slips of the tongue which cannot be described as canonical phoneme exchanges (such as double or overlapping articulations: Laver, 1980) has been a topic of considerable debate. Their low rate of occurrence in transcribed speech errors relative to whole phonemic segment errors (e.g., less than 5% of all sound errors could be accounted for by a single feature error Fromkin, 1971) has prompted researchers to argue against their existence (Shattuck-Hufnagel & Klatt, 1979). More recently, however, detailed analyses of elicited speech errors have demonstrated that phonemes uttered in error differ acoustically from canonical phonemes (Frisch & Wright, 2002; Goldrick &

Blumstein, 2006). Using a tongue-twister paradigm, Goldrick and Blumstein (2006) demonstrated that in an exchange error (*keff* → "geff") the /g/ in "geff" differed in VOT from both a canonical /k/ and a canonical /g/. Similar findings have been reported from articulatory analyses (Boucher, 1994; Goldstein, Pouplier, Chen, Saltzman, & Byrd, 2007; Mowrey & MacKay, 1990; Pouplier, 2003, in press). For example, in an electromyographic (EMG) investigation Mowrey and MacKay (1990) observed transversus/verticalis muscle activity, normally associated with /l/ production, in the production of "bay" during repetitions of *Bob flew by Bligh bay* (Mowrey & MacKay, 1990). These findings not only pose challenges for categorizing responses as 'correct' or 'phoneme exchange error', but also have implications for models of production.

Discrete models of production propose that selection of a phoneme occurs in an all-or-none manner (Levelt et al., 1999). According to these models, the acoustic and articulatory variations described above do not reflect cognitive planning, but instead are attributed to a late stage process of phonetic encoding or motor execution (Levelt et al., 1999). A cascading model of production, however, proposes that acoustic and articulatory variation reflects traces of the partially activated competing phonemes (Goldrick & Blumstein, 2006). Evidence for this account comes from a post-hoc comparison by Goldrick and Blumstein (2006) which revealed that VOTs were affected by lexical status; the VOT in slips resulting in real words (e.g., *kess* → "guess") compared to those resulting in nonwords (*keff* → "geff") tended towards the VOT for a canonical production of that real word. On the other hand, for nonword outcomes the VOT remained more similar to the intended rather than the produced onset. This finding clearly supports a cascading model in which activation at the lexical level results in partial activation at the phonemic level (see Frisch & Wright, 2002, for similar evidence from degree of voicing in /s/–/z/ exchanges). In turn, the existing evidence for cascading suggests that many of the slips of the tongue which have been disregarded in earlier SLIP studies may in fact reflect

partial activations of competitor phonemes.

By focusing on whole-phoneme errors and categorizing responses, existing reports of SLIP experiments may not be capturing important patterns in the data. Responses categorized as 'other errors' may include double, overlapping, or otherwise modified articulations which exhibit clear influences of errorful activation in the speech production system. Moreover, to the extent that articulation can vary continuously, the boundary between 'correct' and 'other' onsets may be as arbitrary as that between 'other' and 'exchange' onsets.

The primary purpose of the present paper is therefore to provide an analysis of elicited slips of the tongue which does not depend on the categorization of responses. We present analyses of spoken responses from a Word Order Competition (WOC) task (Baars & Motley, 1976). As in the SLIP task, participants are required to repeat word pairs in response to a cue. Unlike the SLIP task, this task does not require targets to be primed with repeating onsets. Instead, the tendency to create errorful responses is here manipulated by the use of a 'directional' cue: In a minority of (critical) cases, participants are cued to repeat the words they have seen in reverse order (i.e., they see *tum gop* and are expected to respond "gop tum"). On a proportion of such trials, participants may respond in error (a full phonemic exchange here would yield two real word outcomes "gum" and "top"). Targets are manipulated for potential outcome, referred to as *competitor* below, and are embedded in lists of fillers providing the *context*. Our context manipulation, in a manner similar to Baars et al. (1975) and Hartsuiker et al. (2005), provides the first investigation of how context affects the articulation of speech errors, by crossing (real word or nonword) competitor with (mixed or wholly nonlexical) context.

In order to establish that the WOC task exhibits the lexical bias observed using other paradigms, we first present a transcription analysis of errors as acoustically perceived. We then present an articulation analysis of electropalatographic (EPG) records

from a subset of participants. EPG provides an articulatory record of tongue-to-palate contact over time (Scobbie, Wood, & Wrench, 2004; Wrench, 2003) and allows us to report articulatory patterns that reflect a 'blend' of two phonemes which may not be perceived acoustically (Edwards, Gibbon, & Fourakis, 1997). While EPG has often been used for clinical investigations of pathological speech (Gibbon, 2006), this paper provides the first EPG analysis to our knowledge of slips of the tongue in unimpaired speakers.

To analyze articulatory patterns, we have developed a technique that quantifies spatial and temporal articulatory variation. An advantage of this technique is that it allows us to present an analysis of articulatory patterns without pre-assigning participants' responses to categories (such as 'error', 'correct'). This technique also allows us to evaluate articulation relative to either the intended phoneme or the competitor phoneme, giving an indication of the influence of each. Our experimental design, which manipulates context (whether the participant sees any real words) and competitor (whether a full phonemic substitution would result in a real word or not) allows us to investigate the effects of higher-order speech planning processes on the articulation of individual phonemes.

Method

The experiment took the form of a Word Order Competition (WOC) task (Baars & Motley, 1976). In this task, participants see pairs of words or nonwords, which disappear and are followed by an arrow on the screen. Participants are told to repeat the words they have just seen aloud, in the direction of the arrow (i.e., for *tup golve* followed by a rightward arrow, the correct response is "tup golve"; followed by a leftward arrow, "golve tup"). Items followed by leftward arrows may give rise to onset exchanges rather than full exchanges of words (participants may say "gup tolve" when the prompt following *tup golve* points left). We varied the status of the *competitor*, so that exchanges might result in nonwords, as above, or in real words, such as "gum top" from *tum gop*. We also

manipulated the *context* in which target pairs were embedded, nonlexical (all nonwords) vs. mixed context (nonwords and real words). Each participant saw a full set of target word pairs counterbalanced across two (blocked) context conditions, resulting in a fully within-participants design.

Participants

48 native speakers of English from the University of Edinburgh and Queen Margaret University research community participated in the experiment. Eight of these participants were additionally recorded using electropalatography (EPG). We excluded one of the speakers recorded with EPG from all analyses due to a difficulty with responding to stimuli at the experimental presentation rate. All participants were treated in accordance with the University of Edinburgh and Queen Margaret University ethical guidelines.

Materials

Targets. The targets consisted of 96 CVC(C) nonwords, 24 with each of the onsets (/k/, /g/, /t/, /d/). Each target ended in a bilabial or labio-dental phoneme (/p/, /b/, /f/, /v/, /m/), sometimes preceded by a liquid (/l/). The targets were designed to achieve firm tongue contact with the EPG palate at word onset while minimizing the amount of tongue contact at word offset. Due to restrictions in the targets and in the generation of competitor pairs (see below), only 92 unique nonwords could be used in the experiment. Four of these were repeated to complete the design.

Competitor Pairs. Each velar onset target (/k/, /g/) was paired with an alveolar onset target (/t/, /d/) to generate 48 target pairs, yielding 12 pairs of each onset combination (/d/–/g/, /t/–/g/, /d/–/k/, /t/–/k/). Refer to Appendix A for a complete list of competitor pair targets. Targets were paired on the basis of their *competitors*—that is, the type of outcome that would be observed if participants were to exchange the onsets

of the words in the pair. Half of the pairs had real-word competitors, such that an exchange would result in two phonologically well-formed words; half had nonword competitors, where the result of an exchange would be two nonwords. For each real-word competitor pair, there was a phonologically similar nonword pair: for example, the real-word competitor pair *keam turve* was matched to the nonword competitor pair *keeb turp*. The 24 competitor pairs in each category (real-word and nonword) included six pairs of each of the onset combinations above. For presentation in the experiment, half of each set of six pairs was reversed, so that participants were equally likely to encounter pairs with onsets /t/–/g/ or /g/–/t/.

Foil Pairs. In addition to the competitor pairs, 48 nonword foil pairs were generated consisting of 12 pairs of each combination of /s/–/m/, /s/–/n/, /r/–/m/, and /r/–/n/. These stimuli were generated to obscure the matched alveolar-velar pattern of the competitor pairs. All foil pairs were non-word competitors so that a substitution errors could only yield nonword outcomes.

Contexts. Competitor pairs were embedded in either *nonlexical* contexts (participants never saw real words), or *mixed* contexts (participants saw a total of 62% nonwords and 38% real words, including targets, foil pairs, and context pairs). The two contexts were created independently: the nonlexical context consisted of 150 pairs of nonwords; in the mixed context, 75 pairs of real words were added to a further 75 nonword pairs.

Two lists of competitor pairs were created, each consisting of 12 real-word competitor pairs and 12 nonword competitor pairs, and organized such that if a real-word pair appeared in one list, the corresponding nonword pair appeared in the other list. Each list was combined once with each context, yielding four context-list combinations, comprising 198 pairs each (24 competitor pairs, 24 foil pairs, 150 context pairs).

Competitor and foil pairs were randomly distributed within each list, such that 2-4 context pairs preceded each pair which required a response. Finally, two experimental treatments were created, each consisting of a nonlexical context list and a mixed context list, such that each participant saw every competitor pair over the two lists.

Apparatus

The experiment took place in a sound-treated recording studio at either the University of Edinburgh or Queen Margaret University.

Acoustic–Only Recording. The acoustic signal of participants responses was recorded on to a DAT recorder with a Sony ECM-TS125 condenser microphone and digitally converted into .wav files with a 22,050Hz sampling rate. Stimuli were presented on a 15" LCD monitor using a desktop computer and E-Prime software (Schneider, Eschman, & Zuccolotto, 2002). Audio materials were played over stereo headphones.

Acoustic & EPG Recording. Prior to testing, each participant was fitted with a custom electropalatography (EPG) palate (manufactured by Incidental, Newbury, UK or Grove Orthodontics, Norfolk, UK) molded to fit a dental cast from an impression of the hard palate. The EPG palate is made of acrylic and contains 62 embedded silver contacts on the lingual surface of the artificial palate, organized in seven rows of eight contacts and one row of six contacts. EPG data was recorded at a rate of 100Hz using the WinEPG system (Articulate Instruments Ltd, Edinburgh, UK), which connected the palate to a multiplexer unit that transfered the data to an EPG3 scanner and then to the serial port of a desktop computer. See Figure 1 for an illustration of an EPG record. The acoustic signal of participants' responses was recorded at 22,050Hz using an Audio Technica ATM10a microphone. A desktop computer, to which the microphone and WinEPG system were attached, was used to record participants' responses with Articulate Assistant (Wrench, 2003) software. A laptop computer was used to control stimulus presentation,

using E-Prime (Schneider et al., 2002): participants saw word pairs on a 15" LCD monitor and auditory signals were played over stereo headphones.

Procedure

Once participants had been seated, and fitted with an EPG palate if appropriate, they were instructed that their task was to repeat aloud the word pairs that had appeared on the screen, as quickly as possible and in the direction of an arrow. For example, a correct response to *perch house* followed by a rightward arrow would be "perch house"; when followed by a leftward arrow, it would be "house perch".

All target pairs were followed by a leftward arrow. Foil pairs were also followed by leftward arrows, to prevent participants noticing the alveolar-velar pattern of target pairs. All other items were followed by rightward arrows, creating a 3:1 ratio of rightward to leftward arrows throughout the experiment.

Each word-pair appeared on-screen for 1000ms, after which it was immediately replaced with a (right- or left-ward) arrow. The arrow remained on the screen for up to 1000ms. If a response onset triggered the E-Prime voice cue before the 1000ms deadline, the arrow disappeared. To encourage rapid responses, if a response was not initiated within the 1000ms arrow presentation, a loud buzzer was played together with a red flash on the monitor. The next item pair appeared on the screen 400ms after the response initiation or buzzer warning.

The experiment started with a 10 item practice session containing two mock target items, followed by a break to allow for questions and feedback. The experiment was then presented in two blocks with a one minute break between runs. To keep the method as consistent as possible to similar investigations (e.g., Baars et al., 1975; Hartsuiker et al., 2005) participants listened to brown noise (random noise with a spectral slope of 6dB/octave) throughout the task at a volume as loud as comfortable. The total duration

of the experiment was approximately 15 minutes.

Control Session

The seven participants who were recorded with EPG were also recorded in a control session. The control session was used to collect articulatory data used for comparison in the articulation analysis (see also Articulation Method). We used the same stimulus items as the WOC Task, but only presented the original target pairs, each preceded by one filler pair. The procedure was identical except that all target pairs were followed by a right arrow, cueing the participants to speak the target and filler items in the presented order. For example, in the WOC Task participants were presented with *gope doof* and cued with a left arrow to respond "doof gope", but in this task they were cued with a right arrow to respond "gope doof". The control recording sessions resulted in the presentation of 672 nonword target onsets in which errors were not elicited.

Data Treatment

Following the experiment, we analyzed participants' responses in two ways. The first analysis was based on the perceptual transcription of responses from all 47 speakers. Each target response was assigned to a category, discussed in more detail in *Transcription Method*. The second analysis was based on the articulation data recorded from seven speakers and the method is discussed in more detail in *Articulation Method*. For both analyses, each nonword within every target pair was treated independently: For example, within the target pair *gope doof*, both *gope* and *doof* were coded independently. This is different from many studies using phonological exchange elicitation paradigms, which report outcomes analyzed by pairs of targets (e.g., Baars et al., 1975; Hartsuiker et al., 2005). In part, we adopted this approach because the WOC task encourages the confusion of left- and right-hand targets, making it difficult to distinguish between anticipation (*gope doof* → "doof dope") and perseveration (*gope doof* → "goof gope") errors; in part,

it allowed us to isolate onsets for the articulatory analyses. Also, for both analyses we focus exclusively on target onsets. We use the term *competitor substitution* to refer to slips of the tongue that involve the substitution of the competitor's onset for the target onset.

Transcription Method.

Each target item (e.g., *gope* in *gope doof*) was coded as Correct (*gope* → "gope"), a Competitor Substitution (*gope* → "dope"), or as Other. The 'other' category included all responses that did not fit into the 'competitor substitution' or 'correct' categories. This included cases in which the targets were pronounced in the incorrect order (e.g., the correct response order to gope doof is "doof gope"; an incorrect order would be "gope doof"), which could not be distinguished from a combination of competitor and rime substitutions. Where participants produced more than two (non-)words in response to a target pair, the first two items which constituted a response that resembled the intended target pair were selected as the response.

All of the responses from five speakers (10.6% of the total data) were cross-transcribed by another transcriber to test for inter-rater reliability. There was 96% agreement on categorizing the onset of each target. The transcribers came to a consensus on the transcriptions of the other 4% of items.

Articulation Method.

Articulatory records were created by identifying the cue to speak (or acoustic offset of the previous word) and the onset of the spoken vowel for each item. The EPG contact data was then extracted between these two time points. Each articulatory record was then visually trimmed to include the first palate before full closure through to the first palate after the final full closure release before vowel onset. This trimming method was conducted blind to to the intended articulatory pattern to ensure that a velar closure was equally as likely to be identified as an alveolar closure. Full closure was defined as any

lateral continuous path across the EPG palate. Targets in which velar closure did not yield a continuous path across the palate were reexamined. If closure was achieved across all but the middle two posterior contacts at any point during the articulation, this was treated as full closure and the record was trimmed accordingly. EPG records that did not include this degree of velar closure, or full closure at other positions, were excluded from the articulation analysis because the start and end points could not be reliably identified. This included 10.8% of the target pair items and 9.4% of the control pair items. No more than 14.5% of the data was excluded for any given participant and the excluded items were evenly distributed across cells in the design matrix.

The articulatory analysis did not take into account contact that may have occurred during rime production for two reasons. First, all target pairs were designed to minimize tongue-to-palate contact after onset production in an effort to reduce contact between nonwords within each target pair. Second, while some contact may have been recorded as the result of higher vowels or liquids, the identification of onsets based on the acoustic offset of the previous word minimizes the chance of mistaking rime closure of the first nonword with onset closure of the second nonword of each target pair.

To calculate tongue-to-palate contact variability, we first created reference articulations for each speaker for each relevant place of articulation from the control session EPG record. First we standardized the epochs of all of the trimmed EPG recordings which contained full closure (see Appendix B). Then, we created two average articulations for each speaker: One for velar and one for alveolar articulations.

For each of the relevant EPG recordings from the WOC Task (i.e., all recordings which included full closure at any point) we then calculated difference scores from that speaker's reference articulations. This was done by treating each EPG record as a series of 62-dimensional vectors over time, and measuring the Euclidean distance between experimental and reference articulations once time had been standardized such that there

were equal numbers of epochs for each EPG recording. Appendix B gives more details of the variability calculation we used.

We conducted a difference calculation twice for each EPG recording from the WOC task: once relative to the intended articulation reference, and once relative to the competitor articulation reference. The results of the calculations were measures of deviance in contact from the intended onsets (the higher the score, the less like a 'typical' intended onset a particular recorded articulation was), and of deviance in contact from the competitor onsets (the lower the score, the more like a 'typical' competitor onset). Table 1 give examples of EPG recordings and the derived deviance scores relative to an alveolar and velar reference.

Transcription Results

A total of 4512 targets were presented. Speakers did not respond to 73 target items. Of the 4439 responses recorded, we coded 3003 as correct (67.7%), 51 as competitor substitutions (1.1%), and 1385 as other errors (31.2%). Table 2 shows the numbers of competitor substitutions by condition. To compensate for non-normality due to low competitor substitution error rates, we applied a rau (Rationalised Arcsine Transform Unit: Studebaker, 1985) transformation to the percentage error scores. This is similar to an arcsine transform but retains values close to the original percentage values over most of the range of scores. We report by-participant (F1) and by-item (F2) ANOVA statistics with Competitor (real word, nonword) and Context (mixed, nonlexical) as within-participant and within-item factors and Group (behavioral, EPG) as a between-participant but within-item factor.

The analysis revealed a significant effect of Competitor, showing that competitor substitutions occurred more frequently when the target had a real word competitor (1.6%) compared to a non-word competitor (0.6%; $F1(1, 45) = 10.72$, $p = .002$, $MSE = 38.81$;

$F2(1, 94) = 5.32$, $p = .02$, $MSE = 62.57$). We did not observe significant effects of Context or Group, or any significant interactions.[1]

Transcription Discussion

When the speech produced in the WOC paradigm is transcribed and analyzed in a similar way to previous SLIP studies, there is evidence for a lexical bias, with more competitor substitutions for real word competitors than for nonword competitors. This establishes that the effect of the WOC paradigm on speech production is similar to that of the SLIP task, and adds to the growing body of experimental support for a lexical bias in the substitution of word onsets (e.g., Baars et al., 1975; Hartsuiker et al., 2005; Humphreys, 2002). In the WOC experiment, context has no effect, consistent with Humphreys' (2002) SLIP findings, which have been taken as support for a feedback account of the lexical bias effect (Dell, 1986). According to such an account, lexical bias results from the reinforcement of lexical representations from a backward flow of information from phonological representations. This reinforcement is an automatic process, and occurs independent of context.

Before accepting these findings, however, it is important to note that they are subject to the limitations of transcription and coding. In the present experiment, only 51 responses (1.1% of those recorded) were assigned to the 'competitor substitution' category. A small deviance in the assignment of categories to responses could have had a large impact on the data included for analysis, and on the reported outcome of the experiment. This problem would become particularly acute if tongue slips resulted in articulatory patterns which included properties of both the intended and competitor phonemes. Listeners tend to categorize between-boundary phonemes as belonging to one category (Liberman, Cooper, Shankweiler, & Studdert-Kennedy, 1967; Liberman, 1997; Pisoni & Tash, 1974), a process which is reinforced by the segmental nature of standard

transcription methods (Frisch & Wright, 2002).

A post-hoc inspection of the articulation data recorded from the seven participants in the EPG condition confirmed that 13.2% of recorded responses contained contact in both alveolar and velar palatal regions. It is possible that this pattern reflects activation of both the intended and competitor onset phonemes. Moreover, other articulation patterns revealed atypical alveolar and velar contact (for example, alveolar phonemes such as /t/ or /d/ with unusually posterior contact), which may also have been affected by the context they occurred in or the lexical status of their competitors. Table 1 provides examples of some of the contact patterns recorded; see, in particular, examples (c–d) for 'double articulations' and (b) for less typical alveolar contact.

To take into account the possibility that the 98.9% of responses disregarded in the transcription analysis included relevant slips of the tongue, we conducted a quantitative analysis of tongue contact variability on the EPG data collected during the WOC task. This analysis includes all of the analyzable responses made by seven participants. Unlike the transcription-based analysis, it is not constrained by a segmental transcription system and is not subject to listener biases. Analyzing the EPG recordings in this way abandons the concept of 'types' of tongue slip: Instead, we make the *a priori* assumption that all articulations made under conditions designed to promote speech errors may have been affected in some way. Therefore, we use difference scores between the EPG recordings from the WOC experiment and those collected during the control session. The latter can be defined *a priori* as reference articulations, because they were obtained under conditions that should not promote tongue slips above normal variation.

Articulation Results

We conducted two analyses of the tongue-contact variability data. The first used difference scores calculated relative to the intended onset references; the second used

difference scores relative to the competitor onset references. For both analyses we report by-participant (F1) and by-item (F2) ANOVA statistics with Competitor (real word, nonword) and Context (nonlexical, mixed) as within-subjects factors. See Table 3 for means and standard deviations of variability across conditions.

For the first analysis we used a measure of tongue contact variability calculated relative to the reference articulation of the intended target onset. There was a marginal effect of Context; articulation tended to vary more from the intended target in the nonlexical context than in the mixed context (difference scores of 2.09 vs. 1.93; $F1(1,6) = 4.81$, $p = .071$, $MSE = .004$; $F2(1,47) = 3.44$, $p = .070$, $MSE = .201$). No other effects were significant (all Fs < 1).

The second analysis was based on the tongue contact variability of each target calculated relative to the reference articulation of the competitor onset. This analysis showed a significant effect of Competitor (marginal by-items): Articulation was more similar to the competitor onsets when targets had a real word competitor than when they didn't (difference scores of 4.41 vs. 4.55; $F1(1,6) = 15.71$, $p = .007$, $MSE = .0008$; $F2(1,47) = 3.62$, $p = .063$, $MSE = .226$). No other effects were significant (all Fs < 1.07).

Figure 2 shows how tongue contact varies relative to the intended and competitor reference articulations across the experiment. Points in the upper left corner would be likely to be perceptually transcribed as 'correct', and those in the lower right corner as 'errors'. Importantly, there are also several points representing articulations which are intermediate between the intended and competitor references, raising concerns for any analysis based on the categorisation of responses.

Articulation Discussion

The analysis of articulation variability confirms the pattern of results reported from the transcription analysis; speech is more affected in cases where there is a real word

competitor, but this effect is not sensitive to the nature (mixed, or fully nonlexical) of the context. However, this analysis extends the earlier findings in three important ways. First, it provides evidence that lexical competition affects articulation. Where target onsets have real word competitors (e.g., *gome* could result in "dome"), the articulation is more similar to the competing real word onset (/d/) than in cases where there is no real word competitor (*gofe* → "dofe").

Second, a comparison of the analyses using target and competitor onset references establishes clearly that the differences in articulation must be attributed to the influence of competitor onsets: although *gome* → "dome" results in a 'more /d/-like' onset than *gofe* → "dofe", *gome* → "dome" does not result in a 'less /g/-like' onset than for the equivalent nonword onset competitor. In other words, onsets are attracted *towards* a real word competitor, rather than *away* from the target onset. This observation extends evidence showing that VOT can be affected by lexical status (Goldrick & Blumstein, 2006), because VOT is a continuous measurement, so a 'less /g/-like' /g/ will tend to be 'more /k/-like' along that dimension, regardless of whether the difference is due to the influence of a /k/ (a similar argument can be made for degree of voicing, as in Frisch & Wright, 2002).

Third, this is the first experiment to our knowledge which suggests that there may be an effect of context on articulation: Participants tended to produce articulations which varied more from the intended target onsets when the context consisted entirely of nonwords than when it was mixed. This raises the interesting possibility that context effects in the production of speech errors may have a different locus from competitor effects, in line with models incorporating both feedback and self-monitoring (e.g., Hartsuiker et al., 2005; Nooteboom & Quené, in press). However, the effect of context clearly requires replication before firm conclusions can be drawn. Below, we focus on the more compelling demonstration that articulation is affected by competitor onsets.

General Discussion

This paper reported an experimental investigation of lexical and contextual influences on speech production using the Word Order Competition (WOC) paradigm to elicit slips of the tongue. In the first part of the paper, we transcribed and categorized participants' responses in a manner similar to previous studies. We found that phoneme exchanges were more likely to result in real words than in nonwords, and that this effect was independent of context. This confirmed the general lexical bias found in previous work (e.g., Baars et al., 1975; Humphreys, 2002; Hartsuiker et al., 2005; Nooteboom & Quené, in press), validating the WOC task as an investigative tool. However, like previous speech error investigations, our findings suffered from the fact that the analysis depended on identifying a small subset of the available evidence (here, 1.1% of responses).

The major focus of the present study was an articulatory analysis of every analyzable response made by each of seven participants. In contrast to the analysis of transcribed responses, we made the *a priori* assumption that every utterance a participant made might be affected to a greater or lesser degree by its competitor (what would result if onset phonemes were fully exchanged) and by its context (whether items other than nonwords appeared in the list). Importantly, this analysis revealed that the onsets of targets with real word competitors were articulated more similarly to competing phonemes than were those of targets with nonword competitors. This shows not only that there is a clear lexical bias in articulation, but that it is directly attributable to a real word competitor.

Implications for models of production

Several theorists have proposed that activation in the speech production system can cascade from level to level (e.g., Goldrick, 2006; Peterson & Savoy, 1998; Rapp & Goldrick, 2000). According to such a view, competition, such as that caused by the WOC

paradigm, can result in partial activation of not only the intended but also the competitor onset phonemes. To the extent that the competitor onset is active, it will influence the eventual articulation of a target.

Such a view is clearly compatible with our findings. In cases where there is competition, articulation of the onset becomes more similar to a 'canonical' articulation of the competitor phoneme. Moreover, the clear influence of the competitor is modulated by lexical status: The effect is greatest when a substitution of the competitor onset would result in a real word, whether the context contains real words or not. Importantly, we can rule out the possibility that competition from the lexical level simply adds noise to articulation, because dissimilarity from the target phoneme is not affected by lexical status. Our analysis demonstrates that, for example, variation in /k/ articulation may be more /g/-like or less /k/-like, but not necessarily both.

The most straightforward way of accounting for the overall lexical bias in articulation found in the present study is to incorporate feedback between phonological and lexical representations into the model. If a competitor onset phoneme which would result in a real word receives activation, that activation feeds back from this and the other relevant activated phonemes to the lexical representation. This, in turn, reinforces the activation of the competitor onset. In cases where no real words would result from a slip of the tongue, there are no lexical representations to feed back to, and the competitors remain relatively inactive.

A potential source of caution in interpreting the present results is that the speech analyzed in this paper was obtained using a tongue slip elicitation paradigm, during which participants were auditorily presented with brown noise. These facts might lead one to question whether the variation we report is 'representative'. Essentially, there are two answers to this question. First, observations of acoustically or articulatory deviant speech have been reported using a variety of laboratory methods including repetition (Goldstein

et al., 2007), tongue-twisters (Frisch & Wright, 2002; Goldrick & Blumstein, 2006), and a SLIP task (Pouplier, in press). Similar articulatory variation has been observed where the experiment was not designed to elicit errors (Boucher, 1994). Although listening to noise may increase articulatory variation (the so-called Lombard Effect: Lombard, 1911), in the present study (as in earlier work, e.g., Baars et al., 1975; Hartsuiker et al., 2005) noise was used for the duration of the experiment and cannot account for between-condition differences. Second, even if the systematic differences we report here were not to be found in everyday speech, we believe that they constrain the set of potential models that can be used to account for speech production.

Benefits of measuring variation

The variation analysis was predicated on the assumption that activation could cascade from level to level in the speech production system. This had the consequence that the articulation of any onset recorded under conditions designed to promote speech errors could in theory have been affected to a greater or lesser degree by competitor activation. EPG recordings from the WOC task were indiscriminately assigned *a priori* to the 'error' category and were therefore all included in the analysis. In order to measure the degree to which they deviated from 'normal' articulations, they were compared to averaged onset phonemes obtained under comparable but (by definition) non error-producing circumstances. This approach has a clear advantage over transcription-based methods.

Nooteboom and Quené (in press) have recently noted that studies which rely on categorizations of speech errors regularly suffer from sparse data, rendering the statistical analyses unreliable. This criticism is clearly relevant to the transcription analysis presented here; although we presented the analysis for comparison with previous work, we noted that only 1.1% of transcribed items were entered into the analysis. This is equivalent to only approximately 1 'error' out of 96 responses for each of 47 participants.

By focusing on variation, we were able to include 599 word onsets, equal to approximately 85 out of 96 responses for each of 7 participants.

Using EPG, it would also be possible to assign responses to categories: Indeed, the motivation for several previous instrumental investigations is that responses can not be categorized from the acoustic record (Pouplier & Hardcastle, 2005). However, the problem of categorization remains: Any attempt to delimit classes of response requires a theoretical motivation for positing the classes in the first place. It may of course be the case that the 1385 responses we assigned to the 'other' category during transcription are not affected by the same processes as the 51 competitor substitutions. But if some of them may have been, then the imposed boundaries between categories may result in misleading findings.

Perhaps the most important consequence of analyzing articulatory variation, however, is that this study makes a direct link between the cognitive processes involved in speech production, and the resulting motor movements that produce the speech. Like other newly-emerging paradigms such as mouse tracking (see Spivey, Richardson, & Dale, in press, for a review), articulatory analysis shows that motor movements can give us a fine-grained insight into the cognitive processes that drive them.

Conclusion

In this paper we have demonstrated that variation in articulation reflects higher-level influences on speech production. This observation is important for the ongoing development of speech production models and establishes that measuring articulatory variation can provide an insight into cognitive processing.

References

Baars, B. J., & Motley, M. T. (1976). Spoonerisms as sequencer conflicts: Evidence from artificially elicited errors. *American Journal of Psychology*, *89*(3), 467–484.

Baars, B. J., Motley, M. T., & MacKay, D. G. (1975). Output editing for lexical status in artificially elicited slips of the tongue. *Journal of Verbal Learning and Verbal Behaviour*, *14*, 382–391.

Boucher, V. J. (1994). Alphabet-related biases in psycholinguistic inquiries: Considerations for direct theories of speech production and perception. *Journal of Phonetics*, *22*(1), 1–18.

Byrd, D., Flemming, E., Mueller, C. A., & Tan, C. C. (1995). Using regions and indices in EPG data reduction. *Journal of Speech and Hearing Research*, *38*, 821–827.

del Viso, S., Igoa, J. M., & Garcia-Albea, J. E. (1991). On the autonomy of phonological encoding: Evidence from slips of the tongue in Spanish. *Journal of Psycholinguistic Research*, *20*, 161–185.

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

Dell, G. S., & Reich, P. A. (1981). Stages in sentence production: An analysis of speech error data. *Journal of Verbal Learning and Verbal Behavior*, *20*, 611–629.

Edwards, J., Gibbon, F., & Fourakis, M. (1997). On discrete changes in the acquisition of the alveolar/velar stop consonant contrast. *Language and Speech*, *40*, 203–210.

Frisch, S. A., & Wright, R. (2002). The phonetics of phonological speech errors: An acoustic analysis of slips of the tongue. *Journal of Phonetics*, *30*, 139–162.

Fromkin, V. A. (1971). The non-anomalous nature of anomalous utterances. *Language*, *47*, 27–52.

Garrett, M. F. (1976). Syntactic processes in sentence production. In R. J. Wales & E. Walker (Eds.), *New approaches to language mechanisms* (pp. 231–256).

Amsterdam: North Holland Publishing Company.

Gibbon, F. (2006). *Bibliography of electropalatographic (EPG) studies in English (1957–2006)* (Unpublished Research Report). Edinburgh, UK: Queen Margaret University College.

Goldrick, M. (2006). Limited interaction in speech production: Chronometric, speech error, and neuropsychological evidence. *Language and Cognitive Processes*, *21*(7–8), 817–855.

Goldrick, M., & Blumstein, S. E. (2006). Cascading activation from phonological planning to articulatory processes: Evidence from tongue twisters. *Language and Cognitive Processes*, *21*(6), 649 - 683.

Goldstein, L., Pouplier, M., Chen, L., Saltzman, E., & Byrd, D. (2007). Dynamic action units slip in speech production errors. *Cognition*, *103*, 386–412.

Hardcastle, W. J., Gibbon, F., & Nicolaidis, K. (1991). EPG data reduction methods and their implications for studies of lingual coarticulation. *Journal of Phonetics*, *19*(3), 251–266.

Hartsuiker, R. J., Corley, M., & Martensen, H. (2005). The lexical bias effect is modulated by context, but the standard monitoring account doesn't fly: Related beply to Baars et al. (1975). *Journal of Memory and Language*, *52*, 58–70.

Humphreys, K. R. (2002). *Lexical bias in speech errors.* Unpublished doctoral dissertation, University of Illinois at Urbana-Champaign.

Laver, J. D. M. (1980). Monitoring systems in the neurolinguistic control of speech production. In V. A. Fromkin (Ed.), *Errors in linguistic performance: Slips of the tongue, ear, pen, and hand.* New York: Academic Press.

Levelt, W. J. M. (1983). Monitoring and self-repair in speech. *Cognition*, *14*, 41–104.

Levelt, W. J. M. (1989). *Speaking: From intention to articulation.* Cambridge, MA: MIT Press.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral & Brain Sciences*, *22*, 1–75.

Liberman, A. M. (1997). *Speech: A special code.* Cambridge: MIT Press.

Liberman, A. M., Cooper, F. S., Shankweiler, D. P., & Studdert-Kennedy, M. (1967). Perception of the speech code. *Psychological Review*, *74*, 431–460.

Lombard, E. (1911). Le signe de l'élévation de la voix. *Annales des Maladies de L'oreille et du Larynx*, *37*, 101–119.

Motley, M. T., Camden, C. T., & Baars, B. J. (1982). Covert formulation and editing of anomalies in speech production: Evidence from experimentally elicited slips of the tongue. *Journal of Verbal Learning and Verbal Behavior*, *21*, 578–594.

Mowrey, R. A., & MacKay, I. R. (1990). Phonological primitives: Electromyographic speech error evidence. *Journal of the Acoustical Society of America*, *88*(3), 1299–1312.

Nooteboom, S. G. (2005). Listening to one-self: Monitoring in speech production. In R. Hartsuiker, R. Bastiaanse, A. Postma, & F. Wijnen (Eds.), *Phonological encoding and monitoring in normal and pathological speech.* Hove, UK: Psychology Press.

Nooteboom, S. G., & Quené, H. (in press). Self-monitoring and feedback: A new attempt to find the main cause of lexical bias in phonological speech errors. *Journal of Memory and Language*.

Peterson, R. R., & Savoy, P. (1998). Lexical selection and phonological encoding during language production: Evidence for cascaded processing. *Journal of Experimental Psychology: Learning Memory and Cognition*, *24*, 539–557.

Pisoni, D. B., & Tash, J. (1974). Reaction times to comparisons with and across phonetic categories. *Perception and Psychophysics*, *15*, 285–290.

Pouplier, M. (2003). *Units of phonological coding: Empirical evidence.* Unpublished doctoral dissertation, Yale University.

Pouplier, M. (in press). Tongue kinematics during utterances elicited with the SLIP technique. *Language and Speech*.

Pouplier, M., & Hardcastle, W. (2005). A re-evaluation of the nature of speech errors in normal and disordered speakers. *Phonetica*, *62*, 227–243.

Rapp, B., & Goldrick, M. (2000). Discreteness and interactivity in spoken word production. *Psychological Review*, *107*, 460–499.

Roelofs, A. (2004). Error biases in spoken word planning and monitoring by aphasic and nonaphasic speakers: comment on Rapp and Goldrick (2000). *Psychological Review*, *111*(2), 561–72.

Samuel, A. (1996). Phoneme restoration. *Language and Cognitive Processes*, *11*(6), 647–654.

Schneider, W., Eschman, A., & Zuccolotto, A. (2002). E prime 1.0 [Computer software manual]. Pittsburgh, PA: Psychological Software Tools.

Scobbie, J. M., Wood, S. E., & Wrench, A. A. (2004). Advances in EPG for treatment and research: An illustrative case study. *Clinical Linguistics and Phonetics*, *18*(6–8), 373–389.

Shattuck-Hufnagel, S., & Klatt, D. H. (1979). The limited use of distinctive features and markedness in speech production: Evidence from speech error data. *Journal of Verbal Learning and Verbal Behavior*, *18*, 41–55.

Spivey, M., Richardson, D. C., & Dale, R. (in press). Movements of eye and hand in language and cognition. In E. Morsella & J. Bargh (Eds.), *The psychology of action, vol. 2*. New York: Oxford University Press.

Studebaker, G. A. (1985). A "rationalized" arcsine transform. *Journal of Speech and Hearing Research*, *28*, 455–462.

Wrench, A. (2003). Articulate assistant user guide, version 1.3a [Computer software manual]. `http://www.articulateinstruments.com`.

Appendix A

Target Items used in the WOC Task

Real Word Competitors	Nonword Competitors
gim dulp	gib dulm
gome dasp	gofe dasb
gope doof	gobe doove
dap gime	dalf gipe
dape gam	dabe galf
dulf gamp	duf galve
tave gub	tafe gup
tum gop	tup golve
timp giff	tib gilf
garp tiv	garm tirve
guff tob	gulb tov
gube tolf	goove tolm
keff darve	kem darf
kip doff	kiv dolf
koom darp	coob dalp
dop kuv	dolb kulve
duff cump	dulve culp
dup kive	dulp kife
toop kerm	toove kurp
tove kemp	tofe keb
tome kipe	tobe kive
keam turve	keeb turp
curf talm	kerp talb
kime turb	kibe turp

Appendix B

Quantification Method for EPG

The quantification method used in this paper returns a value, Δ, corresponding the similarity of any two sections of EPG recording (low values of Δ can be interpreted as greater similarity). Δ is calculated in two stages. First, individual sections of EPG are rendered comparable by *standardizing the epoch*; Second, pairwise comparisons between standardized EPG sections are made using *average Euclidean distance.* This combination of procedures ensures that Δ is sensitive to differences in either the timings or the locations of articulatory contact.

The method outlined here differs from more standard EPG analyses (e.g., Byrd, Flemming, Mueller, & Tan, 1995; Hardcastle, Gibbon, & Nicolaidis, 1991) in two respects. First, no emphasis is placed on particular regions of the EPG palate when quantifying articulatory behavior (a more standard approach would be similar to that taken in the 'articulatory analysis' section of this paper); Second, Δ is a ratio variable which is subject to further analysis using standard quantitative methods. This appendix gives a brief definition of the calculation of Δ.

Standardizing the Epoch

EPG consists of a recording over time of contacts with an M-electrode (typically 62-electrode) palate. Here, we treat the palate as an ordered array of contacts, which is sampled periodically (typically every 10ms). A section of EPG recording thus consists of K samples, or K M-element arrays, where K is a consequence of whatever method was used to determine the 'beginning' and 'end' of an articulatory sequence of interest.

In order to render EPG sections comparable, they are converted from records of K samples to records of N (typically 10) epochs, by calculating the amount of contact at

each electrode over each epoch.

At the end of this procedure, all EPG sections for analysis are represented as N-epoch records, where the value of an EPG section x for a given electrode m at epoch n ($x_{m,n}$) is higher the more contact there was at that electrode during the nth part of the articulation (slow articulations with many samples per epoch are likely to result in higher values).

Average Euclidean Distance

Since each record now consists of N epochs of M-electrode EPG, it is straightforward to compare articulations. The difference δ_n between any two articulations at epoch n is the Euclidean distance between relevant vectors. A smaller δ_n implies that the tongue made similar contacts over that epoch of the articulation. To compare two complete articulations, we can calculate mean δ_n over N epochs. For recordings from M electrodes of articulations x and y, quantized into N epochs,

$$\Delta = \frac{\sum_{n=1}^{N} \delta_n}{N} = \frac{\sum_{n=1}^{N} \sqrt{\sum_{m=1}^{M} (x_{m,n} - y_{m,n})^2}}{N}$$

Author Note

Footnotes

[1]For comparison with previous studies we also analysed the untransformed error percentages. This analysis yielded a comparable pattern of results, with only the effect of Competitor reaching significance ($F1(1, 45) = 8.53$, $p = .005$, $MSE = .0006$; $F2(1, 94) = 9.48$, $p = .003$, $MSE = .014$).

Table 1

Sample of EPG recordings of intended alveolars and variability scores from alveolar and velar references: (a) typical alveolar, (b) less typical alveolar indicated by higher distance from alveolar reference, (c, d) double articulations which contain either overlapping (c) or non-overlapping (d) alveolar and velar contact, (e) less typical velar indicated by a higher distance from velar reference, and (f) typical velar articulations.

Token	Reference	
	Alveolar	Velar
a	1.38	5.13
b	3.95	4.59
c	3.15	4.00
d	3.62	3.89
e	4.31	2.56
f	5.17	1.04
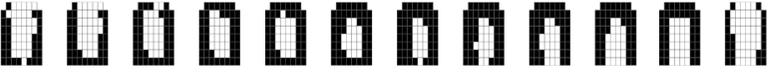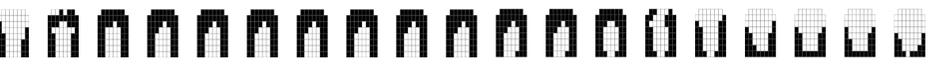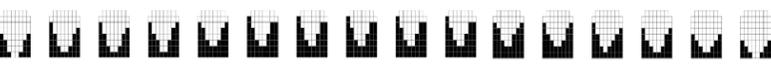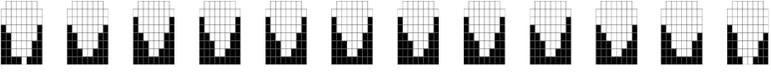
Table 2

Total numbers of competitor substitutions by experimental condition

| Context	Nonlexical	Mixed
Nonword competitor	6	8
Real word competitor	16	21

Table 3

Mean difference scores (arbitrary units) for quantitative analysis of tongue contact variability (SDs in brackets)

Context	Nonlexical	Mixed
Variability from intended onset		
Nonword competitor	2.07 (.20)	1.89 (.17)
Real word competitor	2.12 (.24)	1.98 (.29)
Variability from competitor onset		
Nonword competitor	4.52 (.29)	4.58 (.31)
Real word competitor	4.40 (.35)	4.43 (.27)

Figure Captions

Figure 1. Example of an EPG record

Figure 2. Scatterplot showing calculated tongue contact distance of each recorded onset articulation from the intended and competitor references. Points in the upper left corner represent articulations similar to the intended reference but different from the competitor reference; points in the lower right corner represent articulations similar to the competitor reference but different from the intended reference.

alveolar ridge

left

right

velum

time (10ms samples)