11 Language processing and segmental OCP effects*

Stefan A. Frisch

1 Introduction

The Arabic verbal roots are subject to a long-distance phonotactic constraint that is well known for its implications for autosegmental representation (McCarthy 1986, 1988, 1994). In this constraint, originally proposed as an instantiation of the Obligatory Contour Principle (Goldsmith 1979), repeated place of articulation features are not allowed within a root. Subsequent research has shown that the details of consonant occurrence in the Arabic roots are complex, with the strength of the phonotactic restriction gradiently dependent on the similarity of the consonants involved, the presence of intervening segments, and the contrasts available in the segmental inventory of Arabic (Pierrehumbert 1993; Frisch, Pierrehumbert, and Broe in press).

The gradience of the phonotactic patterns in the Arabic lexicon provide strong evidence for a functional phonetic motivation for the constraint. The similarity avoidance constraint in Arabic is quantitatively dependent on similarity, distance between segments, segment frequency, and segmental position in the word. No formal model that prohibits feature co-occurrence like the autosegmental OCP can capture the richness of the patterning. A wide range of evidence from psycholinguistics suggests that processing a sequence of similar items is more difficult than processing a sequence of dissimilar items. Thus, we can account for the presence of similarity avoidance constraints in the phonotactics of Arabic as a consequence of functional pressure to make language processing as easy as possible. I claim that the richness of phonotactic patterns directly (quantitatively) reflects the functional explanation. In this way, statistical analysis of the lexicon provides a novel type of evidence for functionally motivated constraints and rules out alternative formal explanations (see Hawkins 1994 for similar arguments at the syntactic level). Statistical patterns in the synchronic lexicon arise as the result of a diachronic influence of processing difficulty on language change. Over time, functional pressures on the language have shaped the lexicon by influencing borrowing, the creation of nonce forms, and the loss of lexical items.

Despite the diachronic origin of the similarity avoidance patterns, native speakers are aware of these patterns, so they must be considered a part of the synchronic linguistic knowledge of the speakers. The co-occurrence constraint on homorganic consonants influences metalinguistic judgements of acceptability for novel roots (Berent and Shimron 1997; Frisch and Zawaydeh 2001) and the accommodation of borrowed lexical items (Frisch, Pierrehumbert, and Broe in press).

Moreover, a phonotactic constraint based on processing difficulty should be universal. In fact, similarity avoidance constraints for homorganic consonants like those in Arabic have been found in a wide range of languages, such as English, Javanese, and Ngbaka. Analogous constraints that apply to repeated laryngeal features rather than repeated place features are also attested across unrelated languages such as Sanskrit, Hausa, and Souletin Basque (MacEachern 1999). Further, in cases where lexical patterns have been analysed statistically, the co-occurrence patterns are gradient and quantitatively depend on similarity (Berkley 1995, 2000; Buckley 1997; Frisch 1996; Pierrehumbert 1993).

In this chapter, a functional account of segmental OCP effects is given. Languages avoid sequences of repeated similar segments because they are difficult to serialise. An explanation for the difficulty of repetition can be found in models of language processing. Current language processing models use activation and competition in a neural network of linguistic units to account for similarity-dependent error rates in perception and production. In an activation/competition model of the encoding of a serial sequence, the units (e.g. segments) to be serialised must be activated and deactivated in the proper sequence (Dell, Burger, and Svec 1997). For a segment that has already been encoded, the node in the network corresponding to that segment has fired, and that node must be inhibited so that it does not continue to fire. Simultaneously, for a segment that is soon to be encoded, the node corresponding to that segment must be excited so that it is ready to fire at the proper time. If a sequence involves a repeated segment, the periods of inhibition and excitation may overlap and disrupt encoding of the correct sequence.

The chapter is organised as follows. Section 2 provides an overview of current theories of language processing, focusing on models of segmental processing. Section 3 reviews the segmental OCP patterns in Arabic and other languages, and shows these patterns are gradient and similarity dependent. Section 4 presents the processing account of similarity avoidance constraints and reviews some outstanding problems with this account. Section 5 concludes the chapter with a brief summary.

2 Current theories of language processing

Phonetic and psycholinguistic research on speech perception and production has found that language processing is highly interactive. The perception of structures at the level of speech sounds, spoken words, and larger syntactic constituents involves three general stages: the activation of compatible structures, competition between structures that share common input, and selection of the winning structure (e.g. Swinney 1979; Zwitserlood 1989). In addition to activation and competition within a structural level, there is activation between levels, such that one can influence another (e.g. Ganong 1980). Similarly, the production of speech sounds, words, and phrases appears to involve interactive activation and competition (Dell 1986; Stemberger 1983). Errors in speech production reveal competition between articulatory plans. The processing-based account of similarity avoidance constraints relies on the notion of activation and competition between segments in phonological encoding, so evidence for activation models of phonological processing is presented below.

Theories of phonological processing come from two sources of data. Research on speech perception has sought to determine the organisation and access of phonological information in the mental lexicon. Research on speech production has examined how lexical information is accessed and encoded as a sequence of gestures. In both cases, much of what is learned about language processing comes from studying the systematic errors that are made. Errors tell us something about items that are partially activated during processing. Additional evidence for activation/competition can be gleaned from differences in ease or speed of processing. Tasks involving greater competition between segments or words generally take longer to perform than tasks with no competition.

In this section, I first review evidence in favour of activation/competition from speech perception and speech production. I then introduce connectionism as a formal model of activation/competition in language processing. Finally, psycholinguistic evidence for the processing difficulty involved in repetition is reviewed and connectionist models of serial encoding that explain the processing difficulty are presented.

2.1 Spoken-word recognition and lexical neighbourhood effects

There is a great deal of evidence from speech perception and spoken-word recognition that the phonological lexicon is organised as a multidimensional acoustic-phonetic space, and the recognition of a word takes place when a lexical item becomes sufficiently activated in comparison to its competitors. The organisation of phonological words has been described in terms of groups of words that are phonetically similar, called *lexical neighbourhoods*. For example, the lexical neighbourhood for *cat* would include words such as *bat*, *fat*, *cut*, *kit*, *cap*, *can*, *scat*, and *cattle*. Substantial evidence for lexical neighbourhoods has been found by comparing the performance of experimental participants in processing words that differ in their neighbourhood characteristics (Luce and Pisoni 1998).

Three factors related to the organisation and activation of words in lexical neighbourhoods have been shown to influence the processing of phonological words. It has long been known that more frequent words are easier to process when compared to less frequent words (e.g. Miller, Heise, and Lichten 1951). Presumably, high-frequency words like *buy* are easier to activate than low-frequency words like *bough*. In addition, the *density* of a lexical neighbourhood influences performance. Words that have many lexical neighbours (e.g. *cat*, *lick*) are more difficult to process than words that have fewer lexical neighbours (e.g. *quiz*, *purge*). Finally, the frequency of the neighbours of a target word influences processing. Words with many high-frequency neighbours are more difficult to process than words with mostly low-frequency neighbours. The density of its neighbourhood and frequency of its neighbours determines how much competition the target word receives from other lexical items (Goldinger, Luce, and Pisoni 1989). Words that are high in frequency with only a few low-frequency neighbours are easy to process, while low-frequency words with many high-frequency neighbours are hard to process.

The competition between words and their lexical neighbours influences processing in a variety of tasks (Luce 1986). In an *identification task*, the participant is asked to determine what word has been presented when the word's identity is masked by noise. In this task, ease of processing is reflected in the accuracy in identification, and easy words are identified more accurately than hard words. In a *repetition naming task*, the participant is asked to produce the word as soon as possible after it is presented (either auditorily or visually). In the repetition naming task, ease of processing is reflected in the latency between the presentation of the stimulus and the start of the participant's production of it. Easy words are produced after a shorter delay than hard words. In the lexical *decision task*, the participant is asked to identify whether the stimulus is a word or a nonword. In lexical decision, the influence of lexical neighbourhoods depends on whether the stimulus is a word or nonword. For a response of 'word' for a word (e.g. *kite*), participants are more quickly able to answer if the word is from a dense neighbourhood than a sparse neighbourhood. For a response of 'nonword' for a nonword (e.g. *gite*), participants are more quickly able to answer if the nonword is from a sparse neighbourhood. Both patterns make sense if words are organised into lexical neighbourhoods. A word in a dense neighbourhood will activate a relatively large number of other words, prompting the participant to respond 'word'. For a nonword in a sparse neighbourhood, there will be relatively few words that are activated, so the 'nonword' response will not receive much competition.[1]

Further evidence for the activation of lexical items in processing comes from research on the time course of spoken-word recognition using eye-tracking equipment. In these experiments, participants are given verbal instructions to manipulate objects represented by icons on a computer screen. Before an object can be manipulated, the participant must fixate their gaze upon the object to co-ordinate the movement of the computer cursor to the object. By examining when participants fixate on an object relative to when the name of that object is spoken, it can be demonstrated that participants begin looking at objects before the entire name has been spoken, and further that they are much more likely to fixate on the target object or a lexical neighbour of that object than on a phonologically unrelated object (Allopenna, Magnuson, and Tanenhaus 1998). Assuming that the probability of fixation on an object is a monotonic function of the activation strength of the object's name, the patterns of fixation compare well to a model of phonological processing based on activation and competition. Allopenna et al. (1998) present simulations using the TRACE model (discussed in section 2.3) that provide a very good fit to their eye-tracking data. Overall, we find consistent evidence that the auditory input in speech perception and spoken-word recognition triggers the activation of word and segment units that compete with one another to be recognised as the percept.

2.2 *Phonological encoding in speech production*

The study of speech production has primarily focused on speech errors that occur in spontaneous speech or in error-inducing experiments. In a speech error, some phonological unit is misproduced. Such errors primarily occur from misordering elements in the speech plan, though in some cases no source for the error is apparent. Errors can occur at the level of the phrase, word, morpheme, or segment, with segmental errors being the most common (Fromkin 1971). Examples of segmental speech errors are given in (1).

(1) *frish gotto* for *fish grotto*
blake fruid for *brake fluid*
spicky point for *sticky point*

In speech errors between segments, errors occur most commonly between segments that share many distinctive features. For example, pairs like p/t, r/l, and s/z have a high error rate, while pairs like p/r and s/m have a low error rate. In other words, errors occur more commonly between segments that are similar to one another. Frisch (1996: ch. 4) analysed consonant-error data for English from portions of two naturally occurring speech-error corpora, the Arizona-MIT corpus (Shattuck-Hufnagel 1979) and the Stemberger corpus (Stemberger 1991). Table 11.1 shows the number of errors in each corpus for each level of similarity in the 'observed' column. Similarity was computed using a metric based on shared features and natural classes (Frisch, Pierrehumbert, and Broe in press). The number of errors that would be expected if errors were randomly distributed is given in the 'expected' column. The relative error rate (the observed

Table 11.1. *Consonant segment errors aggregated by similarity for two naturally occurring error corpora*

	MIT-Arizona corpus			Stemberger corpus		
Similarity	Observed	Expected	O/E	Observed	Expected	O/E
0–0.1	72	519.3	0.14	26	197.1	0.13
0.1–0.2	246	416.5	0.59	98	178.0	0.55
0.2–0.3	234	113.5	2.06	100	57.1	1.75
0.3–0.4	195	88.2	2.21	82	40.6	2.02
0.4–0.5	288	93.1	3.09	131	31.8	4.12
0.5–0.6	238	42.4	5.61	82	14.4	5.68

divided by the expected) is given in the 'O/E' column. Table 11.1 shows that error rate between consonants is clearly highly dependent on similarity.

At the word level, errors also tend to occur between words that are similar. For example, in error-inducing experiments, words with the same vowel are more likely to interact in a consonant error (e.g. *bad back* for *mad back*) than words that do not share a vowel (e.g. *bade back* for *made back*, Dell 1984). In segment errors within polysyllabic words, prosodic similarity is also relevant. Errors are more likely to occur between segments in similar word and stress positions (Frisch 2000; Shattuck-Hufnagel 1992). In errors between whole words, similarity also plays a role. Errors of whole-word substitution are more likely for words that share several segments. Examples include the substitution of *recession* for *reception*, *liberal* for *liveable*, and the blending of *correlated* and *corroborated* to produce *corrobolated* (Fromkin 1971).

As in the case of speech perception, the common patterns of speech-error production described above can be accounted for by a processing model based on activation and competition between phonological units. Segments that share features will activate one another via those shared features, increasing the likelihood that a similar, but incorrect, segment will be selected for encoding instead of the intended segment. Analogously, words that share segments will activate one another based on those shared segments, increasing the likelihood that a similar word or part of a word will be selected for encoding instead of the intended item.

2.3 Models

In both perception and production research, connectionist models that employ spreading activation have been most successful at accounting for the effects of activation and competition in language processing in both error-free and error performance (see Elman, Bates, Johnson, Karmiloff-Smith, Parisi, and Plunkett 1996 for an introduction). These models simulate parallel processing over a

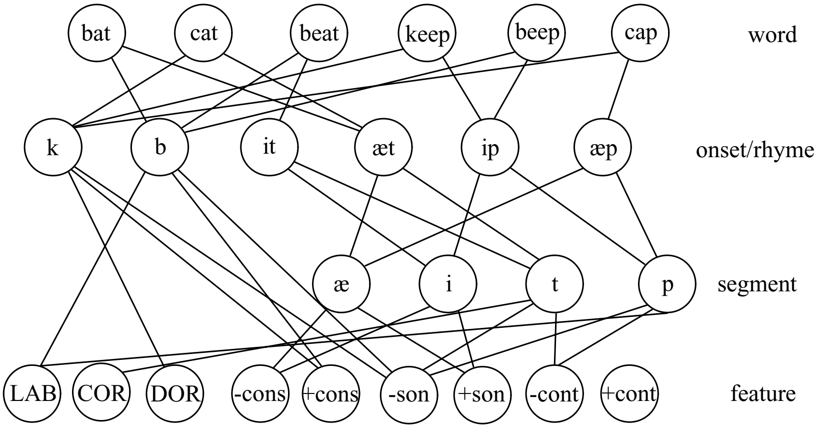

Figure 11.1 Model of phonological encoding, following Dell 1986. Phonological words activate intermediate phonological levels, which in turn activate phonological features to be encoded for articulation.

large set of simple units. Each unit, or node, has an activation level. Nodes are grouped into layers as in figure 11.1. The layers commonly reflect hierarchical layers of linguistic units, such as words, syllables, segments, and features. Each model has an input layer and an output layer that represent the interface between the processing model and semantic, auditory, or articulatory systems. In a model of speech production, words are the input, and the output is a set of segments, features, or articulatory gestures. In a model of speech perception, the input would be acoustic data, distinctive features, or segments, and the output would be a set of units that represent word meanings. In between the input and output are one or more layers of *hidden units*. These layers represent intermediate levels of representation or intermediate stages of processing. In the case of feature-to-word models or word-to-feature models, intermediate layers could include morphemes, syllables, onsets, rhymes, and segments. For example, the TRACE model of spoken-word recognition (McClelland and Elman 1986) has an input feature layer that roughly corresponds to the acoustic features of Jakobson, Fant, and Halle (1952), a word layer for output, and a (hidden) phoneme layer in between.

Processing in a connectionist model is simple. Activation is spread between nodes via weighted connections. In most models, connections are bidirectional, so the units on either end will excite one another when activated, or inhibit one another when not activated. In each time unit of processing, each node's activation level is adjusted according to the amount of excitation and inhibition it is receiving (which is a function of the activation level of nodes connected to it and their connection strengths).

The Dell (1986) model of phonological encoding for speech production, shown in figure 11.1, uses word, syllable, onset, rhyme, segment, and feature nodes. Phonological encoding in this model is simulated by first activating a word that is to be produced (e.g. *keep*). This word spreads activation to corresponding onset and rhyme nodes (e.g. /k/ and /ip/), which in turn activate their corresponding segments (e.g. /i/, and /p/) and features (e.g. [+DOR], [−son]). The onset and rhyme units also activate other words (e.g. *cat* and *beep*) that will introduce competition for the production of the correct segments. The model is run for a number of cycles of spreading activation, after which the most active onset, nucleus, and rhyme segments are selected to be the segments that are encoded. Since the Dell model contains a feature level, it can account for many of the basic facts of segmental speech errors. Similar segments will be activated by shared features, and so will be more likely to be selected for one another, producing an error. In addition, the shared segment nodes between words will cause words with similar segments to be more likely to interfere with one another and result in a speech error.

In some models, the activation of one node inhibits the activation of another node. In the TRACE model of word recognition, words that share segments inhibit one another. Effectively, words in TRACE compete for the segments that are activated by the acoustic feature input. This competition will produce many of the lexical neighbourhood effects discussed above. Words with many neighbours will receive a great deal of inhibition from similar words, and so will be harder to activate over their threshold level. Frequency effects are incorporated in these models by giving units different thresholds of activation for encoding. High-frequency units require less activation to fire than low-frequency units, so high-frequency units are more reliably encoded. Also, high-frequency neighbours to a word will provide more competition than low-frequency neighbours to a word as high-frequency neighbours will be more likely to fire when only partially activated.

2.4 The problem of serial encoding

In spoken-word recognition and speech production, language processing unfolds over time. In processing language over time, it is important not only to recognise or produce the appropriate linguistic units as discussed above, but to also recognise the correct order of elements. The words *pat*, *tap*, and *apt* all involve the same segments, and it is their ordering that differentiates the words. A wide range of psycholinguistic studies have shown that sequences with repeated items are more difficult to process than those that do not contain repetition. Presumably, the problem is that repetition introduces a potential confusion in the linear ordering. For example, when producing *Julius Caesar's seizures*, there is high phonological similarity between *Caesar's* and *seizures*.

All segments but one are repeated, and the two distinct segments /z/ and /ʒ/ are highly similar. The likelihood of a speech error is high, as after planning the sequence /si/, it is potentially unclear whether this is the first /si/ in *Caesar's* or the second /si/ in *seizures*, and the sequences /siz/ and /siʒ/ must both be planned. Consistent with this example, it has been found that sequences containing repeated segments have higher speech-error rates, in both laboratory experiments and naturally occurring errors (e.g. Dell 1984; Stemberger 1983).

It might at first appear that the repetition problem would only exist for production, as there is no serial planning that must take place in perception. However, there is the possibility for perceptual confusion as well in cases where items are repeated. In the case of *Caesar's seizures* the perceptual system must determine that there were two distinct productions of the repeated segments. For example, whatever perceptual mechanism indicates that an /s/ is being heard at the beginning of the first word must be reset before the next word arrives (see MacKay 1970 for a neurologically based account of this phenomenon). If /s/ is activated at the beginning of the next word, is it because there was an /s/ in the previous word, or because a new /s/ has come along? The problem is most acute in cases of immediate repetition of a segment (Boersma 1998). For example, the segmental difference between *heavy oak* and *heavy yoke* is the presence of a /j/ onset on the second word (/i#o/ versus /i#jo/). But, phonetically /i/ and /j/ are basically the same. Unless there is a pause between the words, it is difficult to determine whether /j/-like transitions into the /o/ are the result of coarticulation from a previous /i/ or from a distinct /j/ segment. There is no point during the juncture where the perceptual system can determine that the /i/ has ended and a /j/ begun. This does not imply that it is impossible to perceive a contrast between *heavy oak* and *heavy yoke*. Additional phonetic cues are present, such as the overall duration of the [i/j] and the presence or absence of glottalisation at the word juncture. However, these cues rely on suspending judgement on the identity of a portion of speech until sufficient evidence is accumulated to resolve the uncertainty. So, as in the case of production, there is a certain amount of encoding and sequencing that takes place during perception.

There is substantial evidence that a certain amount of time is required for the perceptual system to detect an item and reset itself to detect the same item again. In visual perception of words, there is a well-known phenomenon of *repetition blindness* (Ericksen and Shutze 1978; Kanwisher 1987). In experiments using very rapid presentation of word sequences, repeated items are sometimes perceptually fused and reported only once, even though there is a brief period of time in between stimulus presentations in which a blank screen is presented. This effect has even been shown for repeated letters within words. For example, in a very rapid presentation of the visual sequence *tell*, *shell*, *oe*, the repeated letter sequence *ell* is often perceived only once, such that participants report seeing *tell*, *shoe* (Morris & Harris 1997).

Analogous perceptual errors on repetition in rapidly presented auditory stimuli have also been reported. Participants listening to rapid sequences of words (that must be artificially sped up to produce the effect) can miss a repeated item, reporting it only once (Miller and MacKay 1994). The same effect has also been found at a more abstract level. MacKay and Miller (1994) played rapid sequences of words in which each phonological word was distinct, with no repetition, but the sequence did contain a sequence of near synonyms. In this case, repetition at the semantic level also produced an increase in misreporting what was heard. Thus, we might conclude that the repetition effect is not entirely auditory or articulatory, but instead results from the need to parse a sequence into discrete units at several levels of abstraction.

Processing difficulties that result from repetition can also be found at the morphological level. Stemberger and MacWhinney (1990) examined morphological errors where speakers fail to mark a past tense verb correctly. For example, the past tense of *kid* is *kidded*, producing a sequence containing repeated /d/ segments. Speech errors in which the past tense was intended but not marked (e.g. *kid* for *kidded* or *walk* for *walked*) were more frequent in the case where the word final segment was phonologically similar to the past tense morpheme. In other words, they found a much higher error rate in cases where past tense marking would produce a repeated sequence of similar segments.

2.5 *Models of phonological encoding*

Many connectionist models of language processing have been adapted from models of reading isolated words, and so do not address the process of serialisation. Serialisation is an integral part of spoken language processing, even for isolated words. As activation-based models of language processing have evolved they have begun to address the problem of serial encoding (e.g. Hartley and Houghton 1996). One solution to the problem is to add a set of *sequencing nodes* that are interconnected with excitatory and/or inhibitory links that would cause them to activate and deactivate in sequence (e.g. Dell, Burger, and Svec 1997). The representation of each word is tied to this cascade of sequencing nodes so that the words' segments (and their corresponding features) would become activated in the correct order. A simplified example of a model of this type is shown in figure 11.2. This model is one example of a model to produce English syllables that is a modification of the Dell 1986 model. The sequencing nodes correspond to the abstract structural units of onset, nucleus, and coda. Activation begins with an abstract word node, which in turn activates the segments of the word and the onset node. Over time, activation of the onset node builds. Activation also begins to build in the nucleus. As activation builds in

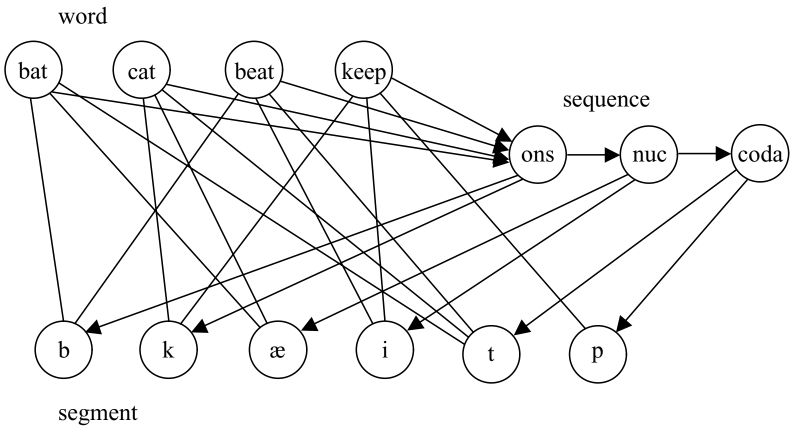

Figure 11.2 Structural model of sequential processing. Independent onset, nucleus, and coda nodes are linked to syllables and their corresponding segments. Activation spreads so that the onset fires first, the nucleus fires second, and the coda fires third.

the nucleus, the activation level of the coda also begins to build. After the onset fires, its activation is lost, and the nucleus is the next most activated element. Once the nucleus fires, activation will continue to build in the coda, until the coda fires. As each sequencing node fires, the most activated segment node is selected to be encoded in that position.[2]

3 The segmental OCP

The phonological phenomena to be accounted for by these language processing models are long-distance similarity avoidance constraints for segments within the lexical patterns of a language. The most well-known of these, the consonant co-occurrence restrictions in the triconsonantal root morphemes of Arabic, were formalised by McCarthy (1986, 1988) using the Obligatory Contour Principle (Goldsmith 1979). The OCP prohibits the repetition of elements on an autosegmental tier. McCarthy proposed, in the case of Arabic, that the OCP is parameterised to operate on privative place of articulation tiers, blocking the generation of a root that contains repeated homorganic segments as C_1C_2, C_2C_3, or C_1C_3. McCarthy also proposed that each place of articulation tier could be selectively sensitive to manner features. These technical refinements to the OCP-Place constraint in Arabic were necessary to account for the fact that roots with C_1C_2 such as those in (2) are not found while those in (3) are relatively rare and those in (4) are quite frequent.

(2) */b m h/
*/t d h/
(3) ?/d s h/
?/t ʃ h/
(4) /d n h/
/t r h/

Segmental OCP effects are found in a wide variety of languages. Place-based consonant co-occurrence constraints are prevalent in the Semitic languages (Bender and Fulass 1978; Buckley 1997; Greenberg 1950; Hayward and Hayward 1989; Koskinen 1964). They have also been found in unrelated languages, including Hawaiian and Serbo-Croatian (MacKay 1970), Ngbaka (Broe 1995), English (Berkley 1995; Frisch 1996), French (Plenat 1996), Italian (Frisch, Pierrehumbert, and Broe in press), Javanese (Mester 1986), Russian (Padgett 1992), and other languages (Yip 1989). Across languages, there appear to be differences in the details of the co-occurrence restrictions for homorganic consonants. For example, in Ngbaka the co-occurrence restrictions are limited to a very few homorganic consonant pairs, and most homorganic consonant combinations are allowed. In Arabic the co-occurrence constraints forbid nearly all homorganic consonant pairs.

When the segmental OCP constraints within a particular language are examined more closely it becomes apparent that homorganic segments are constrained to different degrees, suggesting that the all-or-nothing autosegmental OCP cannot properly explain the patterns. Pierrehumbert (1993) observed that the Arabic consonant co-occurrence constraints depend on the similarity of the homorganic consonant pairs involved. For example, the alveolar consonants divide into distinct series of obstruents and sonorants, where co-occurrence between the series is frequent, while co-occurrence within the series is uncommon (Greenberg 1950; McCarthy 1994). Frisch, Pierrehumbert, and Broe (in press) showed that many additional sub-patterns of stronger and weaker co-occurrence restrictions can be identified for consonant pairs in Arabic. Table 11.2 gives rates of consonant co-occurrence for different levels of consonant similarity for adjacent and nonadjacent Arabic consonant pairs. Similarity is computed using features and natural classes as described in Frisch, Pierrehumbert, and Broe (in press). Nonhomorganic consonant pairs have a similarity value 0, and identical pairs have a value of 1. Observed and expected totals are taken from an on-line version of Wehr's dictionary (Cowan 1979) that contains native and assimilated Arabic triconsonantal roots (Pierrehumbert 1993). Relative occurrence of the combinations (observed divided by expected) is also given, and shows a consistent decrease in co-occurrence rate as similarity increases. There is also a higher rate of co-occurrence for similar but nonadjacent consonant pairs, suggesting that distance between segments also plays a role in the constraint.

Table 11.2. *Consonant co-occurrence between Arabic consonants that are adjacent (C_1C_2 or C_2C_3) and nonadjacent (C_1C_3) in the root by homorganic consonant pair similarity, along with mean wordlikeness ratings on a 1–7 scale from a group of 24 native Arabic speakers*

	Adjacent			Nonadjacent			Wordlikeness
Sim	Obs	Exp	O/E	Obs	Exp	O/E	Rating
0	2978	2349.3	1.27	1411	1248.4	1.13	3.3
0–0.1	451	365.2	1.23	219	203.2	1.08	–
0.1–0.2	492	550.6	0.89	308	283.7	1.09	3.0
0.2–0.3	151	260.2	0.58	96	124.2	0.77	2.9
0.3–0.4	29	131.2	0.22	50	67.0	0.75	2.7
0.4–0.5	14	180.2	0.08	75	103.2	0.73	2.7
0.5–0.6	3	40.8	0.07	8	25.0	0.32	–
0.6–1	0	90.2	0.00	13	40.4	0.32	1.8
thd1	1	199.6	0.01	16	103.9	0.15	2.5T66.09 505.3801

thse1 enralised1bynatvspeakes.1
903(Za)903((201())-(903((coducted1)-(9
254[were12530judged1254[lness2530acc
omandv ordlikveness f(or) 2
similarityfle25(w)w1 of difatence
simoilarity (consonant)-6913(
knoe25(wledgep)-250(of)-250Ar(abic)-250(speak)10(es.1)]TJ19 -1.2 TD[Ion1.((the)-8(1.(o(thre)-8
c-occurr(ence)-8471
OC-Pljac– lfe25(fects–)-371.(in1)2371.(
uted1 OC.s viss

analogous sequences). Coronal forms are robustly attested as in *state*, *stoat*, *stat(istic)*, *stet*, and *astute*. Davis analysed these patterns as the result of a categorical OCP constraint that is active over *sCVC* sequences, where coronals have special status. However, a quantitative analysis of the English lexicon finds that sequences of C_iVC_i are disfavoured more generally (Berkley 1995, 2000; Frisch 1996: Pierrehumbert 1994). It appears that Davis' categorical gap is just an extreme case of the more general pattern of similarity avoidance in English. Frisch (1996; ch. 10) analyses the gap in noncoronal sC_iVC_i sequences as a result of the combined effects of the gradient OCP and a low expected probability of combination of noncoronal sC_iVC_i.

Similar segmental OCP constraints have also been found for laryngeal features (Ito and Mester 1986; MacEachern 1999; Steriade 1982). For example, MacEachern (1999) examined eleven languages that have constraints prohibiting repeated or similar laryngeal specifications within morphemes. As in the case of OCP-Place, the laryngeal co-occurrence constraints across different languages constrain co-occurrence to different degrees. MacEachern showed that the co-occurrence constraints follow a cross-linguistic implicational hierarchy based on similarity. If a language allows two segments with relatively similar laryngeal specifications, such as an ejective and a glottal stop, then the language will also allow two segments with relatively dissimilar laryngeal specifications, such as an ejective and /h/. In earlier work, MacEachern (1997) analysed these languages with a ***SIMILARITY** constraint hierarchy, directly encoding the similarity avoidance constraint. This was later replaced with an analysis using conjoined constraints that ban repeated feature specifications, but the implication that the constraint depends on segmental similarity remains the same.

In the case of Arabic, repeated identical consonants are restricted to the strongest degree of all, and they generally do not occur. For a functional similarity avoidance constraint, identity is the worst possible violation. However, in many languages that have segmental OCP effects, identical segments are treated differently than segments that are highly similar but nonidentical. For example, in Peruvian Aymara, combinations of nonhomorganic ejectives (e.g. /t'/ and /k'/) or homorganic ejectives and aspirates (e.g. /t'/ and /t^h/) are prohibited, but two identical ejectives or two identical aspirates are allowed (MacEachern 1999). In Ngbaka, repeated similar homorganic segments are not allowed (e.g. /p/ and /b/, /b/ and /m/) but repeated identical segments are allowed (Broe 1995). The greatest challenge to the similarity avoidance account is how, in some languages, identity provides a special case that is not considered a violation. The architecture of the language processing module provides a potential solution to this dilemma which is discussed below.

The fact that segmental similarity avoidance constraints appear to be gradient provides important evidence that they are functionally motivated. While many phonological processes are categorical, phonotactic constraints over the lexicon need not be. The lexicon is a large database of phonological patterns. These patterns are potentially influenced by external forces and change over time. It is therefore reasonable to conclude that functional constraints will be directly reflected in phonotactic patterns. In other words, functionally poor patterns will be scarce to the extent that the patterns are poor, and functionally good patterns will be frequent to the extent that the patterns are good. The distribution of patterns in the lexicon can capture the graded nature of phonetic and cognitive functional constraints. These patterns are generally not categorical in their influence on linguistic performance and their impact on the lexicon need not be categorical either. However, this is not meant to imply that the functional patterns are irrelevant to formal linguistic theory. There are a number of cases where the OCP has been used to account for categorical patterns in morphophonological processes (e.g. McCarthy 1988). It is unsatisfactory to consider gradient cases of segmental similarity avoidance as completely unrelated to categorical cases of segmental similarity avoidance. Rather, it is likely that the morphophonological OCP effects are the result of the same functional similarity-avoidance constraint. In these cases the functional constraint has been grammaticised and regularised into a categorical pattern.

A second place where gradient phonological processes can be observed is in sociolinguistic variation. In fact, there is a particularly well-studied case of a variable dissimilation constraint in dialects of English that also sheds some light on the issue of grammaticalisation of functional constraints (Labov 1971; Guy 1991). In a pattern that is analogous to the errors in nonmarking of past tense studied by Stemberger and MacWhinney (1990), varieties of African American Vernacular English have systematic, but variable, nonmarking of past tense in cases where repetition of similar segments would result, as in (5).

(5) [spat] – [spatəd] ‘spotted’
[ɡa^{l}d] – [ɡa^{l}dəd] ‘guided’

Guy (1991) refers to this process as *variable t-deletion*. Patterns of t-deletion are not limited exclusively to past tense morphology or to cases of repetition and the process appears to be phonological rather than morphological. However, contexts where the morphology would introduce repetition show the highest rates of t-deletion, suggesting that the origin of the deletion pattern was in that context and subsequently the pattern has spread more generally through the lexicon. This change in progress could play out in any number of ways, resulting in categorical patterns that only partially reflect the original functional motivation for the change.

There is another phonological phenomenon that, in general, should share common properties with the segmental OCP constraints. These are harmony constraints. Harmony promotes assimilation, rather than dissimilation. As with the similarity avoidance constraints, I am primarily concerned here with non-local harmony. Cases involving assimilation of adjacent segments can be analysed as the result of articulatory simplification, and in these cases, repetition is avoided by reducing a sequence of gestures into a single one (see Jun, this volume). Harmony patterns provide a challenge to a functionally grounded account of similarity avoidance. In addition, there are cases of long-distance consonant assimilation. Any theory that predicts that similarity is problematic for language processing would predict that harmony would be a functionally poor pattern. However, there are many other cases where distinct functional tensions push in opposite directions. The most well known is the tension between economy of articulatory effort and maximisation of acoustic distinctiveness (Lindblom 1990; Jun, this volume; Kirchner, this volume). In the Optimality-Theoretic approach to functional explanation of synchronic patterns, the resolution of these tensions is resolved one way or another through a language-particular constraint ranking. Thus, it may merely be the case that some languages place the functional value of repetition avoidance above the value of harmony, while others do the reverse. This is clearly a fruitful topic for further research (see Kaun, this volume, for further discussion).

4 The language-processing account

The primary conclusion of this chapter, that segmental OCP constraints can be attributed to processing difficulty during serialisation of the segments of a word, has been independently proposed by Berg (1998: 88–98). Berg points out that serialisation is particularly problematic in the case of Arabic, where consonant sequences in root morphemes could easily be confused. Berg and Abd-El-Jawad (1996) present evidence that Arabic consonants are unusually susceptible to speech errors that involve consonant segment reordering, as shown in (6).

(6) /takriib/ for /takbiir/ ‘glorification’
/maraaʕiʃ/ for /maʃaaʕir/ ‘feelings’

Berg claims this error pattern can be attributed to the nonconcatenative morphology of Arabic. Prunet, Béland, and Idrissi (1999) present additional evidence that this is the case. They examined the productions of a bilingual Arabic and French speaker with aphasia. They found he made many misordering errors in Arabic, while those errors were systematically absent in French.

Arabic root structure provides a case where the language is particularly vulnerable to the problem of serialisation during production. Berg argues that segmental OCP effects in Arabic are present because the nonconcatenative morphology makes Arabic particularly vulnerable to the processing difficulty caused by repetition. This vulnerability can be explained by the model of serial encoding presented above. In Arabic, the connections between the root consonants and the sequencing nodes would be weak because the root consonants appear in different syllable positions in the verb form, adjacent to a variety of vowels. Examples for the root /ʃ r b/ are shown in (7).

(7) ʃariba ‘he drank’
ʃurba ‘a drink’
maʃrabun ‘tavern’

Since segmental OCP effects are also found in languages that do not have nonconcatenative morphology, it must be the case that repetition difficulty is a universal of language processing, though the effects of the constraint may be weaker or stronger depending on the other aspects of a language’s phonology and morphology. Similarity avoidance in Arabic is particularly strong due to the special problems of sequencing segments in a nonconcatenative morphological system.

Boersma (1998: 415–40) also takes a phonetic approach to the OCP. Boersma primarily focuses on a perceptual difficulty with repetition, in particular in the case of parsing a sequence of segments where there is immediate repetition. In order to aid recoverability, a language might use epenthesis between repeated consonants. Alternatively, a language might block vowel deletion in cases where repeated identical segments would then result. Boersma considers the perceptual motivation for the OCP to be valid only for adjacent segments. For distant segments, he proposes an articulatory constraint on repetition similar to that of Berg (1998). However, a perceptual motivation for long-distance dissimilation should not be so readily dismissed. In speech perception, a listener must take a rapid stream of air pressure fluctuations and turn it into a sequence of segments, words, phrases, and clauses. Given the high rate of information conveyed in speech, the identification of segments must be very rapid, and consequently may be considered cognitively difficult. A robust speech-perception module has to deal with many cases of missing information or ambiguity (see Wright, Frisch, and Pisoni 1999 for an overview). One possible aid is higher order information, such as the knowledge of the identity of a word being used to supply information about the segments in that word. Thus, the perceptual system does not want to immediately commit to any particular percept of a segment, as additional information might cause a change in the identity of missing or ambiguous segments. If segmental decisions are delayed, or stretched out over time somewhat (i.e. they last longer than the fleeting acoustic pattern corresponding to the segments of the word), then repetition of similar segments within a word may result in blending of perceptual traces and consequently a misperception. For example, the mechanism that recognises segments might still be gathering information about the identity of the first segment in a CVC word when the third segment is spoken. In that case, the new and old consonants would interact and compete, possibly interfering with proper encoding.

Given the preceding discussion, the account of long-distance similarity avoidance constraints is a relatively simple one. Psycholinguistic research and models of language processing suggest that processing repeated items is difficult. In the cases of segmental OCP constraints for place features or laryngeal features discussed above, it appears that repeated gestures using the same articulator are avoided.[3] This is a processing-based explanation, and from a purely phonological standpoint this level of detail is sufficient as a tentative explanation of the prevalence of similarity avoidance constraints in the world's languages.

4.1 Further predictions

Greater insight into the nature of segmental OCP effects can be gained by considering whether additional phonological predictions can be made by a processing-based account. Activation-based models of lexical access and phonological encoding provide a good basis for an account of the repetition constraint. In particular, activation and competition in these models is sensitive to the similarity of phonological units, so segments will compete with one another to the extent that they are similar. Whether the processing model uses some type of activation cascade or is based on a recurrent network, repeating similar items will create competition between the repeated segments, thus disrupting serialisation.

There is a second aspect of these models of serialisation that has not, so far, been discussed. As the serial processing of a word progresses, elements that have been successfully encoded have a greater and greater influence on the encoding of the remaining segments. For example, in the *cohort model* of spoken-word recognition, the candidate set of competing words narrows as phonetic input is processed, until finally a unique candidate is activated (Marslen-Wilson and Welsch 1978). For example, upon hearing [k^h], many words are possible (e.g. *cat*, *cut*, *cute*, *kite*, *kit*). Once most of the segments in a word have been perceived, it is very likely that lexical access will converge on the correct word, and very unlikely that any segment will become highly activated unless it is a possible completion of the word. For example, after [k^{h}ep] there are only a few possibilities (*capes*, *caped*, *caper*, *capon*). A similar context effect has been demonstrated in speech production. Sevald and Dell (1994) showed that it is more difficult to say sequences of words in which initial segments are repeated (e.g. *cat cab*) than it is to say words in which final segments are repeated and initial segments are not repeated (e.g. *cub tub*). This is another example of the repetition effect. For *cat cab*, the repeated word-initial segments increase the amount of competition between words, making it more difficult to encode the correct final segment (/t/ or /b/). However, in cases where the repeated segments are word final (*cub tub*) there is no evidence of

Table 11.3. *Onset and coda consonant co-occurrence in English monomorphemes as a function of homorganic consonant similarity. Data are given separately for each of the first three syllables*

	Syllable 1			Syllable 2			Syllable 3		
Sim	Obs	Exp	O/E	Obs	Exp	O/E	Obs	Exp	O/E
0	1258	1099.6	1.14	558	491.2	1.14	101	99.4	1.02
0–0.1	145	143.2	1.01	64	66.2	0.97	13	17.2	0.76
0.1–0.2	317	275.1	1.15	227	213.7	1.06	67	49.4	1.36
0.2–0.3	162	150.5	1.08	146	117.1	1.25	39	30.3	1.29
0.3–0.4	98	112.6	0.87	61	71.9	0.85	16	17.0	0.94
0.4–0.5	42	73.6	0.57	29	44.4	0.65	14	20.2	0.69
0.5–0.6	56	92.8	0.60	65	79.9	0.81	20	26.3	0.76
0.6–1	14	35.7	0.39	24	31.4	0.76	11	8.2	1.34

lexical competition, and the activation of repeated segments actually facilitates encoding.

Since preceding context helps to ensure correct encoding of later segments in the word, it should be the case that the influence of the segmental OCP constraint is sensitive to preceding context. In other words, the negative effects of repetition on processing should be ameliorated somewhat for repeated segments that occur nearer to the end of lexical items. This hypothesis has been investigated for English and Arabic OCP-Place constraints (Frisch 1996, 2000). A comparison of the strength of OCP-Place constraints for Arabic C_1C_2 and C_2C_3 consonant pairs found some evidence that the C_2C_3 constraints are weaker (Frisch 2000). In English, where the overall OCP-Place constraint is not as strong as in Arabic, there is also evidence that the constraint is weaker nearer to the end of lexical items. Table 11.3 shows relative rates of co-occurrence (O/E) for similar homorganic onset and coda consonant pairs in CVC syllables of English monomorphemes for each of the first three syllable positions in the word. Similarity of the consonant pairs is determined using features and natural classes (Frisch, Pierrehumbert, and Broe in press). Identical consonant pairs are not included, as the details of identical consonant co-occurrence in English are complex (see Frisch 1996 for discussion). While the data for later syllables in the words are somewhat sparse, there is a clear trend towards greater co-occurrence. There is higher O/E for OCP-Place violations away from the beginning of the word, especially in the third syllable. In the third syllable, there is no clear evidence of a segmental OCP constraint.

There is one further sub-regularity in the OCP-Place constraints of Arabic, and it too provides evidence that quantitative patterns in the lexicon reflect the subtle influence of functional processing constraints. Note that the analysis of the segmental OCP as a similarity avoidance constraint predicts that

Table 11.4. *Influence of consonant probability on Arabic onset and coda consonant co-occurrence*

	Consonant probability	
Sim	Low	High
0	1.30	1.25
0–0.1	1.17	1.26
0.1–0.2	0.56	1.08
0.2–0.3	0.26	0.72
0.3–0.4	0.14	0.25
0.4–0.5	0.05	0.08
0.5–0.6	0.00	0.09
0.6–1	0.00	0.00
1	0.00	0.01

all consonant combinations of relatively equal similarity should be relatively equally represented as violations of the constraint. In fact, however, the distribution of violations of the OCP within a similarity group is unequal, and appears to be influenced by consonant frequency. Table 11.4 shows consonant co-occurrence for violations of the OCP-Place constraint in Arabic divided into high-frequency and low-frequency consonant groups. Relative to their expected frequency (O/E), high-frequency consonant pairs are more likely to appear as OCP-Place violations than low-frequency consonant pairs. Since the O/E measure takes expected frequency into account, it appears that high-frequency consonant pairs are less constrained by the OCP than low-frequency consonant pairs. This sub-regularity can be accounted for by lexical processing models that use activation/competition.

Recall that lexical neighbourhoods are an important influence on lexical processing. Also note that high-frequency consonants are found in words in dense lexical neighbourhoods, and low-frequency consonants are found in words in sparse neighbourhoods. In this case, the frequency-based difference in O/E for equivalent OCP-Place violations can be seen as a lexical neighbourhood effect, or equivalently as the influence of particular word exemplars on the constraint. If a particular consonant pair is attested as a violation of OCP-Place, it is likely that the violation can serve as a template for other words to violate the constraint using the same consonant pair. Since high-frequency consonant pairs are more frequent than low-frequency consonant pairs, it is more likely that a high-frequency violation will be found in the lexicon. If this violation then serves as an analogical model, the effect will be one where the ‘rich get richer’, as the violation is supported by its lexical neighbours. In other words, already high-frequency consonant pairs gain an additional advantage due to neighbourhood density that will make them more robust violators of OCP-Place than low-frequency consonant pairs.

As in the case of the gradient nature of the constraint itself, the further gradient effects of position-in-word and consonant frequency on the segmental OCP provide additional evidence that the constraint is functionally motivated. If the OCP constraint were a purely formal feature co-occurrence restriction, these processing factors should have no effect on the constraint, and gradient interactions such as these should not be found.

4.2 The reduplication problem

There is one other unresolved challenge to a functionally based segmental OCP constraint. Languages sometimes employ reduplication as a morphological process. Reduplication creates sequences of repeated segments. It could be that reduplication shares some of the benefits of vowel harmony, as discussed by Kaun (this volume). However, reduplication has been a long-standing problem for connectionist models of language processing, so it would be worth considering more carefully how reduplication and processing interact. Using standard models and training procedures, neural networks cannot learn to generalise the copying process in reduplication to words outside of the training set. This contrasts with other morphological processes like affixation and nonconcatenative morphology, where generalisation to novel words is automatic (Gasser 1998). The lack of automatic generalisation for reduplication in connectionist models appears to be a substantial failure in their ability to provide a general processing model for reduplication. One solution to processing reduplication is to implement a special module in the model that performs the copying when it is needed. This module is similar to the recurrent layer of a recurrent network, providing a memory for the items to be reduplicated. Thus, in a language with reduplication, it may be that special processes are developed to circumvent the repetition problem for reduplicative morphemes.[4]

Though this solution might appear ad hoc, it actually sheds light upon an outstanding problem in the cross-linguistic pattern of similarity avoidance constraints. Recall that some languages appear to have an exception to the laryngeal or place-based OCP constraint for identical segments. The functional difficulty in processing identity, and only identity, can be avoided via a recurrent repetition structure of the same sort that is used for reduplication. While repeated identical segments can benefit from a repetition node, similar but nonidentical segments can never be aided in this manner, as they are not true repetitions. Whether the dissimilarity constraint is based on gestural or perceptual difficulty, there is no way to implement a special copying process for nonidentical segments, so highly similar but nonidentical segments will always be functionally bad. Since maximally similar but nonidentical segments are always the most restricted, the special needs of the processing model offers an account of a cross-linguistically true generalisation about long-distance segmental OCP effects that has not previously been explained (cf. MacEachern 1999). Treating certain OCP problems in parallel with reduplication has also been proposed in formal analyses of the OCP (see Gafos 1998 and Rose 2000 for discussion).

5 Conclusion

Segmental similarity avoidance constraints, like the OCP-Place constraint in Arabic, can be functionally motivated using current theories of language processing that predict difficulty in processing repetition. These constraints have been shown to be gradient, such that lexical co-occurrence patterns reflect degrees of well-formedness with respect to the functional constraint. The gradient nature of these constraints provides strong evidence for their functional motivation. For similarity avoidance constraints, the degree of co-occurrence is a function of similarity, both within individual languages and in typological patterns across languages (Frisch, Pierrehumbert, and Broe in press; MacEachern 1999). Similarity plays a key role in the functional motivation of the constraint. To the extent that repeated items are similar, the repeated items will be difficult to individually activate and encode in the proper serial sequence in perception and production. Similarity leads to mutual activation and competition that interferes with correct identification and serialisation. To the extent that activation and serialisation are universal properties of language processing, segmental OCP constraints can be explained through universal forces that shape all languages.

References

Allopenna, P. D., J. S. Magnuson, and M. K. Tanenhaus (1998). Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory and Language* 38: 419–39.

Bender, M. L. and H. Fulass (1978). *Amharic verb morphology*. East Lansing, MI: Michigan State University African Studies Center.

Berent, I. and J. Shimron (1997). The representation of Hebrew words: Evidence from the obligatory contour principle. *Cognition* 64: 39–72.

Berg, T. (1998). *Linguistic structure and change: An explanation from language processing*. Oxford: Clarendon Press.

Berg, T. and H. Abd-El-Jawad (1996). The unfolding of suprasegmental representations: A cross-linguistic perspective. *JL* 32: 291–324.

Berkley, D. M. (1995). Variability in obligatory contour principle effects. *CLS* 30: 1–12.

(2000). *Gradient OCP effects*. PhD dissertation, Northwestern University.

Boersma, P. (1998). *Functional phonology: Formalizing the interactions between articulatory and perceptual drives*. The Hague: Holland Academic Graphics.

Broe, M. (1995). Specification theory and Ngbaka cooccurrence constraints. Paper presented at the Mid-Continental Workshop on Phonology, Columbus, OH.

Buckley, E. (1997). Tigrinya root consonants and the OCP. *Penn Working Papers in Linguistics* 4: 19–51.

Cowan, J. (1979). *Hans Wehr: A dictionary of modern written Arabic*. Wiesbaden: Otto Harrasowitz.

Davis, S. (1984). *Topics in syllable geometry*. New York: Garland.

Dell, G. S. (1984). Representation of serial order in speech: Evidence from the repeated phoneme effect in speech errors. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 10: 222–33.

(1986). A spreading activation theory of retrieval in sentence production. *Psychological Review* 93: 283–321.

Dell, G. S., L. K. Burger, and W. R. Svec (1997). Language production and serial order: A functional analysis and a model. *Psychological Review* 104: 123–47.

Dell, G. S., C. Juliano, and A. Govindjee (1993). Structure and content in language production: A theory of frame constraints in phonological speech errors. *Cognitive Science* 17: 149–95.

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

Elman, J. L., E. A. Bates, M. H. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.

Eriksen, C. and D. Schutze (1978). Temporal factors in visual information processing: A tutorial review. In J. Requin (ed.), *Attention and performance VII*. Hillsdale, NJ: Lawrence Erlbaum.

Frisch, S. A. (1996). *Similarity and frequency in phonology*. PhD dissertation, Northwestern University.

(2000). Temporally organized representations as phonological units. In M. Broe and J. B. Pierrehumbert (eds.), *Papers in laboratory phonology V: Language acquisition and the lexicon*. Cambridge: Cambridge University Press, 283–98.

Frisch, S. A., J. B. Pierrehumbert, and M. Broe 'in press'. Similarity avoidance and the OCP. *NLLT*.

Frisch, S. A. and B. A. Zawaydeh (2001). The psychological reality of OCP-Place in Arabic. *Lg* 77: 91–106.

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

Gafos, D. (1998). Eliminating long-distance consonant spreading. *NLLT* 16: 223–78.

Ganong, W. F. (1980). Phonetic categorization in auditory word perception. *Journal of Experimental Psychology* 6: 110–25.

Gasser, M. (1998). Representation, generalization, and modularity in connectionist models of language. Paper presented at the Program in Linguistics Colloquium, University of Michigan.

Goldinger, S. D., P. A. Luce, and D. B. Pisoni (1989). Priming lexical neighbors of spoken words: Effects of competition and inhibition. *Journal of Memory and Language* 28: 501–18.

Goldsmith, J. (1979). *Autosegmental phonology*. New York: Garland.

Greenberg, J. (1950). The patterning of root morphemes in Semitic. *Word* 5: 162–81.

Guy, G. R. (1991). Explanation in variable phonology: An exponential model of morphological constraints. *Language Variation and Change* 3: 1–22.

Hartley, T. and G. Houghton (1996). A linguistically constrained model of short-term memory for nonwords. *Journal of Memory and Language* 35: 1–31.

Hawkins, J. A. (1994). *A performance theory of order and constituency*. Cambridge: Cambridge University Press.

Hayward, K. and R. Hayward (1989). 'Guttural': Arguments for a new distinctive feature. *Transactions of the Philological Society* 87: 179–93.

Ito, J. and R. A. Mester (1986). The phonology of voicing in Japanese. *LI* 17: 49–73.

Jakobson, R., G. Fant, and M. Halle (1952). *Preliminaries to speech analysis*. Cambridge, MA: MIT Press.

Kanwisher, N. (1987). Repetition blindness: Type recognition without token individuation. *Cognition* 27: 117–43.

Koskinen, K. (1964). Kompatibilität in den dreikonsonantigen hebräischen Wurzeln. *Zeitschrift der Deutschen Morgenlandischen Gesellschaft* 144: 16–58.

Labov, W. (1971). Variation in language. In C. Reed (ed.), *The learning of language*. New York: Appleton-Century Crofts, 187–221.

Lindblom, B. (1990). Explaining phonetic variation: A sketch of the H and H theory. In W. Hardcastle and A. Marchal (eds.), *Speech production and speech modeling*. Dordrecht: Kluwer, 403–39.

Luce, P. A. (1986). *Neighborhoods of spoken words in the mental lexicon*. (Research in spoken language processing technical report 6) Bloomington, IN: Speech Research Laboratory, Indiana University.

Luce, P. A. and D. B. Pisoni (1998). Recognizing spoken words: The neighborhood activation model. *Ear and Hearing* 19: 1–36.

MacEachern, M. R. (1997). Laryngeal similarity effects in Quechua and Aymara. *WCCFL* 15: 337–51.

(1999). *Laryngeal cooccurrence restrictions*. New York: Garland.

MacKay, D. G. (1970). Phoneme repetition in the structure of languages. *Language and Speech* 13: 199–213.

MacKay, D. G. and M. D. Miller (1994). Semantic blindness: Repeated concepts are difficult to encode and recall under time pressure. *Psychological Science* 5: 52–5.

Marslen-Wilson, W. and A. Welsch (1978). Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology* 10: 29–63.

McCarthy, J. J. (1986). OCP effects: Gemination and antigemination. *LI* 17: 207–63.

(1988). Feature geometry and dependency: A review. *Phonetica* 43: 84–108.

(1994). The phonetics and phonology of Semitic pharyngeals. In P. Keating (ed.), *Papers in laboratory phonology III: Phonological structure and phonetic form*. Cambridge: Cambridge University Press, 191–283.

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

Mester, R. A. (1986). *Studies in tier structure*. PhD dissertation, University of Massachusetts, Amherst.

Miller, G. A., G. A. Heise, and W. Lichten (1951). The intelligibility of speech as a function of the context of the test materials. *Journal of Experimental Psychology* 16: 329–55.

Miller, M. D. and D. G. MacKay (1994). Repetition deafness: Repeated words in computer-compressed speech are difficult to encode and recall. *Psychological Science* 5: 47–51.

Morris, A. L. and C. L. Harris (1997). A sublexical locus for repetition blindness: Evidence from illusory words. In M. G. Shafto and P. Langley (eds.), *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum, 519–24.

Padgett, J. (1992). OCP subsidiary features. *NELS* 22, 335–46.

Pierrehumbert, J. B. (1993). Dissimilarity in the Arabic verbal roots. *NELS* 23: 367–81.

(1994). Syllable structure and word structure. In P. Keating (ed.), *Papers in laboratory phonology III: Phonological structure and phonetic form*. Cambridge: Cambridge University Press, 168–90.

Plenat, M. (1996). De l'interaction des contraints: Une étude de case. In J. Durand and B. Laks (eds.), *Current trends in phonology: Models and methods*. Salford, UK: University of Salford Press, 585–617.

Prunet, J.-F., R. Béland, and A. Idrissi (1999). The mental representation of Semitic words. Ms., University of Quebec, Montreal.

Rose, S. (2000). Rethinking geminates, long-distance geminates and the OCP. *LI* 31: 85–122.

Sevald, C. and G. S. Dell (1994). The sequential cueing effect in speech production. *Cognition* 53: 91–127.

Shattuck-Hufnagel, S. (1979). Speech errors as evidence for a serial order mechanism in sentence production. In W. Cooper and E. Walker (eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett*. Hillsdale, NJ: Lawrence Erlbaum.

(1992). The role of word structure in segmental serial ordering. *Cognition* 42: 213–59.

Stemberger, J. P. (1983). *Speech errors and theoretical phonology: A review*. Bloomington, IN: Indiana University Linguistics Club.

(1991). Apparent anti-frequency effects in language production: The addition bias and phonological underspecification. *Journal of Memory and Language* 20: 161–85.

Stemberger, J. P. and B. MacWhinney (1990). Form-oriented inflectional errors in language processing. *Cognitive Psychology* 18: 329–54.

Steriade, D. (1982). *Greek prosodies and the nature of syllabification*. PhD dissertation, MIT.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)considerations of context effects. *Journal of Verbal Learning and Verbal Behavior* 18: 645–59.

Wright, R., S. Frisch, and D. B. Pisoni (1999). Speech perception. In J. G. Webster (ed.), *Wiley encyclopedia of electrical and electronics engineering*. New York: John Wiley & Sons, 175–95.

Yip, M. (1989). Feature geometry and cooccurrence restrictions. *Phonology* 6: 349–74.

Zawaydeh, B. A. (1999). *The phonetics and phonology of gutturals in Arabic*. PhD dissertation, Indiana University.

Zwitserlood, P. (1989). The locus of the effects of sentential-semantic context in spoken-word processing. *Cognition* 32: 25–64.

Notes

[*] I would like to thank Gary Dell, Bruce Hayes, Janet Pierrehumbert, Dan Silverman, and especially Robert Kirchner for their comments on this chapter. Parts of this work were presented at Johns Hopkins University, the University of North Carolina, and the 75th meeting of the Linguistic Society of America and I would like to thank those audiences for their comments. This work was supported in part by the *Language Learning* Research Assistant Professorship at the University of Michigan.

1. Note that in this case the presentation of a word from a dense neighborhood prompts a faster response than the presentation of a word from a sparse neighborhood, so it is not always the case that greater lexical density makes the psycholinguistic task more difficult. Rather, greater lexical density always leads to greater activation/competition between words, and the effect of that activation depends on the task.
2. Another solution to processing activation patterns over time is to add a special layer of nodes that represent a memory for the context preceding the current state of the network. This layer is called a *recurrent layer*, and is usually implemented as a layer that is identical in structure to the hidden layer. The recurrent layer records the activation level of each node of the hidden layer from the previous time step. This allows the model to have some knowledge of its previous state which can then be used to guide the activation pattern in the hidden layer to the next state (e.g. Elman 1990). A recurrent network such as this could be used in the task of encoding words (e.g. Dell, Juliano, and Govindjee 1993). Informally speaking, the processor knows what sequence of segments is trying to be produced from the input word, and the recurrent layer reflects how far in the sequence of internal steps of processing it has proceeded. This knowledge is then used to change to the appropriate internal state and produce the next segment in the sequence.
3. The gestural account of Arabic co-occurrence is not entirely suitable (McCarthy 1994) and perceptual factors may also be involved (Zawaydeh 1999).
4. In cases where reduplication does not preserve identity, the reduplication process results in a less grievous violation of similarity avoidance. In these cases, the functional difficulty is at least partially circumvented. To my knowledge, connectionist models of language processing have yet to be applied to cases of non-identity in reduplication.