

Evidence for perceptual hypercorrection in American r-dissimilation: A pilot study

Nancy Hall, Bianca Godinez, Megan Walsh, Irene Orellana & Coleen Villegas*

Abstract. We tested Ohala’s (1993) theory that dissimilation results from perceptual hypercorrection for assimilation. We created nonce words by splicing syllables containing /r/ to continuations that either did or did not contain another /r/. When listeners were asked to type these nonce words, they were significantly more likely to omit the first /r/ if there was a later /r/. This is consistent with Ohala’s claim that one rhotic can perceptually mask the presence of another rhotic. The patterns of r-dropping mimics the characteristics of real English r-dissimilation, in which speakers tend to drop the first /r/ from words like *surprise*. We argue that perceptual errors are a plausible cause of the actuation of r-dissimilation, although other articulatory or processing constraints may contribute to its persistence.

Keywords. dissimilation; rhotics; speech perception; language change

1. Introduction. Dissimilation is the avoidance of identical or similar sounds within a certain domain or degree of proximity. A classic historical example is Grassman’s Law, which de-aspirated the first of two aspirated stops in Greek and Sanskrit (as in Sanskrit *bhabhūva* > *babhūva* ‘became’). While Grassman’s Law was apparently a regular process, it is common for dissimilation to have a stochastic character. For example, dissimilation of liquids is common in the history of Romance languages, as in Latin *arbore* > Spanish *arbol* ‘tree’, but was not completely predictable. Dissimilation sometimes becomes grammaticized in the form of morphophonemic alternations or morpheme structure constraints, but here we will focus more on the ‘actuation’ of the process, i.e. the point at which words diachronically develop a new representation.

Dissimilation is in several ways a mysterious process. First, it is unclear why the proximity of similar sounds should be avoided. Languages more commonly display the opposite pattern, assimilation, in which sounds change to become more similar to those near them. Ohala (1993:249) claims that “we are unable to invoke any principles of speech production that would predict change in this direction.” Second, it is unclear why dissimilation is disproportionately attested for certain features (such as aspiration, nasality, laterality, rhoticity and rounding) while being rarely if ever found for some other features (such as voicing or continuancy). Finally, it is unclear why dissimilation seems to be less regular than other phonological changes. Even in languages where its effects are fairly pervasive, there is usually some uncertainty as to whether any given word will undergo dissimilation. For these reasons, the cause of dissimilation is a long-standing problem. The Neogrammarians first raised many of the ideas that are still debated today, from speech errors (e.g. Paul 1880, versus Meringer & Mayer 1895), to articulatory difficulties, to ‘psychological causes’ (Osthoff & Brugman 1878) that we might now express as issues with serialization or other processing constraints.

1.1. DISSIMILATION AS PERCEPTUAL HYPERCORRECTION. In a pair of seminal papers, Ohala (1981, 1993) argued that the unusual characteristics of dissimilation are explained if the process

* California State University Long Beach (CSULB). This work was supported by the National Institute of General Medical Sciences of the National Institutes of Health under Award Numbers UL1GM118979, TL4GM118980, and RL5GM118978. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We also gratefully acknowledge the support of CSULB’s Office of Research and Sponsored Programs.

originates from perceptual hypercorrection for phonetic assimilation. Ohala suggests that the phonological features that dissimilate are precisely those whose phonetic realizations are relatively spread-out, and that temporally extended coarticulation creates ambiguity as to how many instances of a feature are present.

It is known that liquids, for example, can affect vowel formants across several syllables. This type of coarticulation has been most extensively studied in British English, where Tunley 1999 found that F2 and F3 were lowered two syllables before and after an /r/, and raised over the same span near an /l/. Kelly & Local (1986) dubbed these long-range effects ‘resonances’ (see also Heid & Hawkins 2000, West 1999a). Furthermore, listeners are sensitive to these long-range effects. West (1999b) found that when a sentence had one portion excised and replaced with white noise, such as *I said [white noise] today*, listeners could choose with better than chance accuracy whether the removed word had contained an /r/ or /l/ (e.g., *berry* versus *belly*), apparently based on their perception of coarticulation on vowels not adjacent to the liquid. Heinrich et al. (2010) and Tunley (1999) found that synthesized speech is more comprehensible, especially in noisy conditions, if the long-range effects of /r/ and /l/ are added.

However, Ohala points out that drawn-out phonetic realizations can make it hard for listeners to determine how many rhotics are present in a word. The long-range effects of one sound can perceptually mask the presence of a similar sound elsewhere. In a word with two rhotics, like *surprise* [səˈpraɪz], the rhoticity of the first vowel may be misanalyzed as an effect of coarticulation with the [ɹ] in the second syllable. If a listener factors out this assumed coarticulation, they will think the word’s intended representation was /səˈpraɪz/. The listener will then store the word with only one rhotic, and produce it that way in careful speech. In this way, perceptual hypercorrection can lead to dissimilation.

Since listeners are generally good at perceptual correction, such errors will not happen consistently, accounting for dissimilation’s typical lexically irregular character. Listeners may also make occasional errors in the other direction, misinterpreting r-coarticulation for additional rhotics. Such errors of ‘hypocorrection’ would lead to assimilatory insertion of rhoticity, as sometimes occur in pronunciations like [fəˈmɪljə] for *familiar* and [ʃəbət] for *sherbet*.

1.2. CHALLENGES FOR THE HYPERCORRECTION ACCOUNT. For decades, Ohala’s hypercorrection account was widely cited, but empirically untested. As Garrett & Johnson (2013) note, “while there are a number of laboratory demonstrations of correction [for coarticulation], there are almost no controlled observations suggesting that listeners hypercorrect in speech perception...this may be a gap in the literature, but it is an important one.” Since then, there have been several attempts to produce perceptual dissimilation in laboratory settings, but none have succeeded in producing hypercorrection that mimics historical patterns of dissimilation. This has led to spreading skepticism as to whether dissimilation is caused by misperception.

As far as we are aware, the first attempt to produce perceptual dissimilation in the laboratory was undertaken by Abrego-Collier (2013). She tested English listeners’ categorization of a synthetic /ɹ...l/ continuum in words that included either another [ɹ], another [l], or neither. The stimuli had shapes like [ˈaXaYa], where X the ambiguous liquid, and Y is a conditioning segment, either [ɹ], [l], or [d]. The prediction is that X should be more likely to be heard as [r] if Y is [d] or [l], and that X should be more likely to be heard as [l] if Y is [d] or [ɹ]. This experiment essentially attempts to reproduce the kind of r → l and l → r dissimilation that has often happened in the history of Romance (*arbore* > Spanish *arbol* ‘tree’; *colonel* > Middle French *coronel*).

The results did not confirm the hypothesis. Conditioning [ɪ] had no effect on the categorization of X. Conditioning [l] had some effect, but in the opposite of the predicted direction: X was more, not less, likely to be categorized as [l] in a word containing a later [l]. Apparently listeners hypocorrected, producing long-range assimilation. This was the opposite of Ohala's prediction.

Harrington et al. 2016 found a similar lack of evidence for perceptual rounding dissimilation in Italian. Historically, in the development of Latin to Italian, words containing two rounded consonants lost rounding on the first, as in **kwinkwe* > *kinkwe* > [tʃinkwe] 'five'. Since few words in modern Italian contain multiple [w]s, Harrington et al. attempted to produce perceptual dissimilation across word boundaries. They tested a synthesized /kw...k/ continuum, in the words *canto* ... *quanto*, inserted into carrier phrases where the following words either contained /w/ (*quattro*) or did not contain /w/ (*sette*). The prediction was that the ambiguous word should be heard more often as *canto* when it preceded *quattro* than when it preceded *sette*. Again, the prediction was not met: there was no significant difference in [k] vs. [kw] responses depending on the following word.

Ironically, perhaps the only study that has found long-range perceptual hypercorrection was looking for its opposite. Ozburn (2016) explored whether the typologically common pattern of sibilant harmony (requiring identical place of articulation for multiple sibilants in a word) could be an effect of perceptual hypocorrection, as also suggested by Ohala (1993). Using synthesized sVCV-[-VCV] continua, she found that the ambiguous sibilant was more likely to be categorized as [s] in a [__aʃa] context than in a neutral [__ama] context. In other words, she found long-range perceptual dissimilation, with listeners tending to hear [saʃa] rather than [ʃaʃa]. Yet sibilants, cross-linguistically, tend to assimilate rather than dissimilate in place. So, while Ozburn's study suggests that perceptual hypercorrection can occur under some circumstances, it casts further doubt on whether such perceptual errors are related to actual sound changes.

1.3. LIMITATIONS OF PREVIOUS HYPERCORRECTION STUDIES. Abrego-Collier (2013), Harrington et al. (2016) and Ozburn (2016) have greatly moved forward the debate over Ohala's theory of sound change, by demonstrating that it is not trivial to produce long-range perceptual hypercorrection or hypocorrection in the laboratory in ways that mimic historical processes of long-range dissimilation or harmony. However, one issue with all three experiments was that each attempted to produce a type of assimilation or dissimilation that does not currently exist in the language studied.

While phonetic coarticulation has certain universal tendencies, many details of its implementation are language-specific. For example, many languages are described as nasalizing vowels near nasal consonants. Yet there may be differences in exactly how strong the nasal coarticulation is, how far it extends in each direction, how it interacts with vowel quality and length, how it is blocked or not by prosodic or morphological boundaries, etc. Furthermore, the question of whether coarticulation creates perceptual ambiguity depends partly on what other phonological contrasts the language has and how they are phonetically realized. For example, coarticulatory vowel nasalization has more potential to create ambiguity in a language that has phonemic vowel nasalization—but the level of ambiguity would also depend on whether that language had redundant phonetic cues for phonemic nasalization, such as an effect of nasalization on vowel height. For these reasons, we would not expect that every pattern of perceptual hypercorrection can be produced in every language. Rather, each language likely has only a few points of phonetic vulnerability to hypercorrection, and an experiment will only succeed in eliciting perceptual errors if it targets the kind of sound structures that are ripe for such errors.

For example, Abrego-Collier's experiment attempted to produce dissimilatory $r \rightarrow l$ changes in intervocalic English liquids, such as [aɹaɹa] \rightarrow [alaɹa]. While this is a common type of liquid dissimilation in Romance (*arbore* > *arbol*), English has never been prone to such changes. American English does have an active process of r-dissimilation (discussed below), but it differs from the Latin process in two crucial respects. First, it targets either unstressed [ə] or [ɹ] adjacent to consonants, never intervocalic [ɹ] as in [aɹaɹa]. Second, it works by either changing [ə] to [ə], as in *surprise* [sə'praɪz], or by deleting [ɹ], as in *entrepreneur* [ˌɑntɹəpə'nɪ]. It never changes [ɹ] to [l], or vice versa (Hall 2009).

In all likelihood, the fact that some Romance languages dissimilate /r/ to /l/, while English does not, relates to level of phonetic similarity between rhotics and laterals in each language. As Posner 1961:101 notes, in the history of Romance "confusion between l-sounds and r-sounds is frequent even when there is no dissimilatory or assimilatory influence at work". English speakers, on the other hand, do not tend to confuse [ɹ] and [l]. In short, there is little reason to think that American English has the necessary phonetic preconditions for $r \rightarrow l$ or $l \rightarrow r$ perceptual dissimilation, and it is therefore not surprising that Abrego-Collier's experiment could not find it.

Ozburn's (2016) attempt to produce long-range sibilant harmony has the same limitation: it was conducted in English, which has never had sibilant harmony. While, as Ozburn points out, this has the benefit that speakers do not have a grammaticized rule of sibilant harmony that might affect their judgments, it also leaves open the likelihood that languages prone to sibilant harmony may have a different phonetic realization of sibilants than English does.

The experiment that comes closest to testing a language-appropriate process is that of Harrington et al. (2016), which attempted to reproduce in Italian a type of rounding dissimilation that occurred in the history of Italian. However, the process in question is an ancient one; *cinque* for *quinque* was already attested in Vulgar Latin (Posner 1961:78). It may well have been rooted in phonetic coarticulation patterns that no longer exist. Furthermore, the historical process apparently applied only within the domain of a word, while the attempt at modern replication had a word boundary between the target and trigger /w/ (*quanto quattro*). Even if the phonetic pattern of rounding coarticulation were identical in old and contemporary Italian, it could well be bounded to the domain of a word, in which case perceptual dissimilation across word boundaries would not be expected.

In short, an important precondition for testing whether perceptual hypercorrection can produce dissimilation is to use a language that has the right phonetic assimilation patterns to produce the hypercorrection in question.

1.4. R-DISSIMILATION IN AMERICAN ENGLISH. American English has long had a sporadic but pervasive pattern of r-dissimilation. Although Hempl (1893) was the first linguist to describe it, spelling patterns suggest that it has been active since at least the 1700's (Stephenson 1956:275). Hall (2009) lists over 100 reported examples. Throughout this discussion, /r/ is used to refer to both [ə] and [ɹ].

American r-dissimilation most often involves changing an unstressed [ə] to [ə], or deleting an unstressed [ɹ], in a word containing another [ə] or [ɹ]. Usually the /r/ that drops is the first /r/ in the word. Typical examples include *surprise* [sə'praɪz], *governor* ['gəvənə], and *caterpillar* ['kætə,pɪlə]. Such pronunciations are not obligatory, but are very common, unstigmatized, and generally go unnoticed. The domain of the process is flexible: the two /r/s may be in adjacent syllables, or separated by as many as two intervening syllables, as in *thermometer* [θə'mamətə] or *particular* [pə'tɪkjələ]. There are a few strict limitations on the position of the dissimilating /r/: it is never word-initial or word-final, nor clearly intervocalic (there are just a few examples of

dissimilation in [əV] sequences, and none in clear-cut [VɹV] sequences). There do not seem to be any segments that block dissimilation when they occur between the /r/s. While in Latin, l-dissimilation is blocked by an intervening /r/, American r-dissimilation is not blocked by an intervening [l], as seen by *particular*.

While the large majority of cases follow the basic pattern above, there are exceptions. In rare cases, [ə] in prevocalic position changes into another vowel, as in *peripheral* [pə'ɹfiəl] or *barbicate* [bɑːbɪtʃuət] (both apparently influenced by spelling). These two words are also unusual in that /r/ that drops is the second /r/ in the word.

Most /r/s that drop are unstressed, but some speakers also drop stressed coda [ɹ] in words like *quarter* ['kwɔːə]. There are a very few words where [ɹ] drops from stressed onset clusters, such as *frustrated* ['fʌstɹeɪɪd] and *library* ['laɪˌbɹɪ], but such pronunciations are stigmatized.

Hall 2009 argues that typical American r-dissimilation has several characteristics consistent with a perceptual hypercorrection account. The [ə]s that dissimilate are in positions where they are least perceptible, being unstressed and flanked by consonants. The sound they dissimilate to, [ə], is essentially what is left of [əɹ] if a listener perceptually factors out the rhoticity. It is also interesting that the words with the highest dissimilation rate tend to have [əɹ] flanked by one coronal and one labial consonant. Since American [əɹ] has coronal and labial components itself, the local environment may further contribute to the perceptual masking of rhoticity.

These characteristics, plus the fact that the process is active in English today, make this process suitable for perceptual testing. In the experiment described below, we test whether listeners' perception of /r/ is affected by the presence of another /r/ in the same word.

2. Methods

2.1. STIMULI. To test whether one /r/ can perceptually mask the presence of another /r/, we compared the perception of a syllable containing /r/ when spliced to two different continuations: one containing a second /r/ and one containing no /r/s. Data collection was approved by the CSULB IRB (project # 18-217).

Test stimuli consisted of 26 matched pairs of nonce words, such as [mɑː'ɹɪkjulə] vs. [mɑː'ɹɪkjuləm], as illustrated in Figure 1.




Figure 1: Construction of spliced stimulus pairs

Each stimulus was spliced together from portions of two natural speech recordings (produced by the second author). One portion, the ‘target’, contained an /r/ that we believed was most likely to drop. The target portion was identical (the same recording) in both members of the stimulus pair. The other portion contained either a second /r/, as in [-kjulə], or no additional /r/, as in [-kjuləm]. In each pair, the control continuation had the same syllable count, stress pattern, and as far as possible the same segments as the trigger continuation, except that of course /r/ was replaced by another sound. The two continuations were both cut from different recordings than the target portion, so that each stimulus pair uses pieces of three recordings. The three source recordings

were spoken with highly consistent speech rate and intonation, so that the resulting spliced stimuli sound quite natural and the splices are inaudible.

The nonce words were intended to sound like possible English words, and to mimic the prosodic structure of words that are actually prone to dissimilation in English. To achieve this, we began by evaluating Hall’s (2009) list of real English words that have been reported to undergo r-dissimilation. Three student judges eliminated those words that they had not heard with dissimilation in our region (Southern California). The remaining words were sorted by syllable count and stress pattern. We identified the following patterns as common, and created nonce words to follow the same patterns (Table 1).

'σ σ σ <i>comfo(r)ter</i> <i>gove(r)nor</i> <i>ape(r)ture</i>	'σ σ,σ σ <i>cate(r)pillar</i> <i>resp(ir)atory</i> <i>adve(r)sary</i>	σ'σ σ σ <i>pa(r)ticular</i> <i>ve(r)nacular</i> <i>the(r)mometer</i>	σ'σ <i>su(r)prise</i> <i>be(r)serk</i> <i>Be(r)nard</i>
'pæmfə-dʒə 'dʌnfə-rə	'mædvə, tɛ.ɪ 'sænə, toɪ 'nɛpə, tonə	kɑɪ' tɪfə-ɪ ɛn' tɪəmə-ɪ mɑɪ' nɪkjulə	hə' voɪ pə' maɪd sə' biɪ θə' moɪt tə' mət
'σ σ,σ <i>repe(r)toire</i> <i>rese(r)voir</i> <i>afte(r)wards</i>	,σ σ'σ σ <i>co(r)poration</i> <i>hype(r)baric</i>	σ'σ σ <i>pe(r)formance</i> <i>f(r)ustration</i> <i>pe(r)verted</i>	other <i>approp(r)iate</i> <i>furthe(r)more</i> <i>cereb(r)al</i> <i>barbitu(r)ate</i>
'faɪndə, bɑɪt 'kɪnsə, twɑɪ	,mɒrə' dɔɪfən ,mæntɪə' noɪ	dʒə' toɪmɪn sɑɪ' fɛsə soɪ' vərəd vɑɪ' lənə voɪ' tɪeɪfən	'pɪəmə, dæn 'weɪdə, nɛt fə' ɪbɪ ɛk' soɪdiəm

Table 1: Typical prosodic patterns of English words undergoing r-dissimilation, with real examples and nonce stimuli

In 22 of the 26 words, the target portion was first, followed by the trigger or control portion. In the other 4, the target portion was after the trigger/control portion. This is similar to the distribution of dissimilating /r/s in English, where it is most often the first /r/ that drops. The full list of stimulus pairs is given in Table 3.

An additional 7 pairs of nonce words were intended to be test items. However, after completion of data collection, we discovered that there had been errors in splicing: in some, the splice location was not between the two /r/s, while in others, the target portion was not consistent. These items were excluded from analysis.

To mimic as much as possible the experience of hearing a new word in conversation, nonce words were embedded in a recorded English sentence and associated with a picture. For example, [kɑɪ' tɪfə-ɪ] sounds like it could be the name of a medication, so it was presented in the sentence “The doctor has me take [kɑɪ' tɪfə-ɪ]” and associated with a picture of a generic medicine

bottle. Other examples are given below. The nonce word was always in sentence-final position (Table 2).

Sentence	Picture
Pass me the ['kɪnsə, twaɪ].	unidentifiable gadget
He seems kind of [, məntɪə'noɪ].	man resting head on laptop
Did you get the [dʒə'toɪmɪn]?	unidentifiable gadget
This is Mrs. ['dʌnfə-rə].	cartoon image of a woman

Table 2: Sample frame sentences and image description

First /r/ target			Second /r/ target		
Target	Trigger	Control	Trigger	Control	Target
'pæmfəd	- ʒə	- ʒɪn	'pɪam-	'pam-	ə, dæn
'faɪndə	- baɪt	- bɪt	'wɛɪ-	'wɛn-	də, net
'kɪnsə	- , twaɪ	- , twan	fə'ɪ-	fən'ɪ-	bɪl
'mædvə	- , tɛɪ	- , tɛni	ɛk'soɪ-	ɛk'so-	diə-m
, mɒrə	- 'doɪʃən	- 'dɒʃən			
'nɛpə	- , tonə	- , tonəs			
'sænə	- , toɪ	- , toni			
dʒə	- 'toɪmɪn	- 'tomin			
sɑɪ	- 'fɛsə	- 'fɛso			
soɪ	- 'və-rəd	- 'vɛntəd			
vɑɪ	- 'lənə	- lɛnəm			
voɪ	- 'tɛɪʃən	- 'tɛɪʃən			
kɑɪ'tɪ	- fəɪ	- fəzɪ			
ɛn'tɪə	- məɪ	- məkɪ			
mɑɪ'nɪ	- kʒulə	- kʒuləm			
, məntɪə	- 'noɪl	- 'noml			
hə	- 'voɪ	- 'von			
pə	- 'maɪd	- 'mænd			
sə	- 'bɪɪ	- 'bid			
θə	- 'moɪt	- 'mant			
tə	- 'mæt	- 'mɛt			
'dʌnfə	- rə	- rɛn			

Table 3: Test and control items

The audio recordings and pictures were presented to listeners via videos created through Powerpoint. There were two counterbalanced versions of the video. Each version contained 53 slides: 3 practice slides, 16 or 17 slides with two /r/s 16 or 17 slides with one /r/ (these numbers include the 7 pairs later excluded) and 17 filler slides that contained no /r/s. Each slide was timed to be viewed for 10 seconds.

2.2. PARTICIPANTS AND PROCEDURE. Twenty CSULB students participated in this study, ranging in age from 18 to 31 (12 female, 6 male, 1 transgender male, 1 non-binary). All participants reported that they were native speakers of American English, were not regularly exposed to other languages before the age of 5, and had no history of hearing, language or speech disorders. They were compensated with gift cards.

Ten participants viewed each version of the video. Sitting in a quiet room, they were instructed to listen to each audio recording carefully as it played on a computer. Using another computer, they were to type the final word of each phrase in English spelling, writing it the way it sounded. This procedure tended to result in a variety of different spellings for each nonce word, as illustrated in Table 4.

Nonce word	Participant spellings					
'pæmfə-dʒə	<i>pamferger</i>	<i>pampherger</i>	<i>panferger</i>	<i>pamferger</i>	<i>pampherger</i>	
	<i>pampherjer</i>	<i>pamferger</i>	<i>pimpferger</i>	<i>Pampherger</i>	<i>pampherdger</i>	
kaɪ'tɪfəzəl	<i>cartiphosol</i>	<i>cartiphisol</i>	<i>cartiferzol</i>	<i>catifuzole</i>	<i>cartifisal</i>	
	<i>cartiphisol</i>	<i>cartifisol</i>	<i>cartifizol</i>	<i>cartiphozol</i>	<i>Cartiphisol</i>	

Table 4: sample typed responses

2.3. RESULTS. Multiple raters independently coded the typed responses, indicating those where an /r/ present in the audio was missing from the spelling. The r-less spellings are shown in Table 5. For completeness, we have included cases where the trigger /r/s were dropped from spelling, but we cannot tell whether these are dissimilatory, since the trigger portions were not tested without another /r/ present.

Target /r/ drop before trigger /r/	Responses	Target /r/ drop in control condition	Responses
'sænə'toɪi	<i>sanitori, sanitorri</i>	dʒə'tomɪn	<i>Jutomun</i>
dʒə'toɪmɪn	<i>dutormin</i>	kaɪ'tɪfəzəl	<i>catifuzole</i>
soɪ'və-rəd	<i>sevirted</i>	ɛk'sodiəm	<i>Exodium, Excodium</i>
vɑɪ'lə-nə	<i>valerner</i>		
maɪ'nɪkjulə	<i>moniculer, monicular</i>	Trigger /r/ drop	
sə'bɪɪ	<i>subir</i>	ɛn'tɪamə-l	<i>intramual, interominal</i>
tə'mət	<i>tummert, temert</i>	'weɪdə,net	<i>weathernet, weidernet, wadernet</i>
fə'ɪbəl	<i>ferriblawl, feribical</i>		
'weɪdə,net	<i>Wheredanet</i>		

Table 5: responses with missing *r*

Target /r/ was more likely to drop when spliced to a continuation containing another /r/, compared to the control condition, as shown in Table 6. The difference in drop rates between conditions reaches significance in a chi-square test without Yates' correction ($\chi^2 = 4.9$, $p = .027$). We conclude that the presence of a second /r/ does increase the likelihood of listeners perceptually missing an /r/.

	target /r/ not written	target /r/ written
control condition	4	256
trigger /r/ condition	13	247

Table 6: target /r/-drop counts by condition

The patterns of /r/-dropping seen in participants' written responses generally resemble the characteristics of /r/-dissimilation in real American English words. In 11 of the 13 drops, the spelling is consistent with replacement of [ə] by [ɹ], or deletion of coda [ɹ]. The only exceptions are the two responses to [fə'ɪbɹl] (*ferriblawl*, *feribical*), where the [ɹ] seems to be replaced by [l] and [ɪk], respectively. This was also the only test word in which a target [ɹ] in second position underwent dissimilation. The introduction of [k] is surprising on a phonological level, but may be influenced by the large number of English words ending in the morpheme [-ɪkəl] (*theoretical*, *tyrannical*, *heretical*, etc.) While this would indicate some non-phonetic influence on the result, it is still notable that the *-ical* morpheme was heard only in a dissimilatory context.

There were also some cases where listeners inserted extra /r/s in the spelling, as shown in Table 7. While the experiment was not designed to test r-insertion, we note that such insertions are a prediction of the perceptual account of dissimilation: if listeners are uncertain how many /r/s are present in a word, they can err in either direction.

Nonce words	Spellings with extra r
,mɒɹə'dɔɪʃən	<i>mordortorshin</i> , <i>Mortardorshin</i>
kɑɹ'tɪfəzɹl	<i>cartiferzɹl</i>
,mæntɪə'nɒml	<i>manternormal</i>

Table 7: r-insertion in spellings

3. Discussion: Other theories of dissimilation. The results of this experiment are consistent with the hypercorrection account of dissimilation: listeners were less likely to type a letter *r* in a word containing two /r/ sounds than in a word containing only one. However, we acknowledge that this is not the only possible interpretation of the results. Several other functional grounds for dissimilation have been proposed, and could possibly influence performance in the task described here. In this section we describe and discuss these alternative approaches.

The main alternative to the perceptual theory of dissimilation is the idea that dissimilation is driven by factors related to motor planning. These may be either purely physical and articulatory, or related to broader cognitive processing constraints on serialization of repeated elements.

The idea that repeated articulations are difficult to process and produce is quite old; for example, Carnoy (1918:104) claims that dissimilation occurs “when two sounds or two syllables coincide and have to be visualized together and articulated together. In that case the image of one of them easily crowds out the image of the other, and both speakers and hearers hardly realize that one of the repeated members has been omitted...such eliminations occur more often when a physical difficulty is added to the mental strain”. Meringer (1908:91-92) points out that speech

errors sometimes have a dissimilatory character, with words like *unglaublich* ‘incredible’ produced as *ungraublich*.

Modern work has lent some support to this idea. Psycholinguistic literature shows higher speech error rates for sequences containing repeated elements, as well as ‘repetition blindness’ in both visual and auditory stimuli (see Frisch 2004:354-5 for an overview). Frisch 2004 argues that neural network models of language processing predict problems with serializing identical elements: repeated elements require the same node to be activated, inhibited and reactivated in quick succession, and the inhibition period of the first element may overlap with the activation period of the second.

Walter 2007 argues that repetition of articulatory gestures is also physically effortful. She finds that vowels flanked by consonants of identical place are longer than those flanked by consonants of different place, suggesting that speakers have difficulty repeating similar gestures close together.

Whether driven by articulatory or processing constraints, there is evidence that speakers to some extent avoid repeated sounds deliberately. Corpus studies find that speakers avoid the proximity of similar sounds in word formation and word juxtaposition. For example, Martin 2007:82 shows that the presence of identical liquids is statistically under-represented in English neologisms, baby names, brand names and names coined for fantasy role-playing games. Mondorf 2003 finds that morphological choices may avoid nearby identical liquids: if an adjective ends in /r/, this increases the chance that it will take an analytic comparative form (*more bare*) rather than a synthetic comparative (*barer*), where two /r/s would be closer together.

The theories mentioned above do not exhaust the list of approaches to dissimilation; there are approaches that incorporate multiple factors. For example, Gallagher (2010) argues that co-occurrence restrictions between laryngeals, including but not limited to dissimilatory restrictions, are grounded in a desire to maximize perceptual contrastiveness between lexical items. Laryngeal contrasts are relatively poorly perceived in the vicinity of other laryngeals, so allowing multiple laryngeal contrasts in proximity leads to confusability. On this account, there is still a link between perception and dissimilation, but perceptual *errors* are not a necessary mechanism. Rather, dissimilatory restrictions are implemented as OT constraints in the framework of the Dispersion Theory of Contrast. Many other approaches simply instantiate dissimilation in the grammar through rules or constraints, with little attention to the functional grounding of the constraints or the problem of how dissimilation is diachronically activated.

In short, dissimilation has been variously attributed to the difficulty of (a) perceiving, (b) processing, or (c) producing repeated elements. Within each of these three broad functional approaches, there is also debate about the extent to which the avoidance or removal of repeated elements is ‘teleological’, meaning actively sought on some level, or purely accidental (i.e., produced through errors).

3.1. RELATING OUR RESULTS TO ALTERNATIVE EXPLANATIONS OF DISSIMILATION. The task in our experiment (listening to and silently writing nonce words) involves no obvious articulatory effort. Although listeners may to some extent mentally rehearse pronunciations when writing, it seems unlikely that motor constraints could in any way directly lead them to not write /r/.

Mental serialization constraints are a more plausible alternative explanation. However, it is not clear that serialization constraints can explain why dissimilation happens over the distances seen in real American words (*particular* and *thermometer*), or in the nonce stimulus [mɑːˈnɪkjʊlə], which yielded the spellings *moniculer*, *monicular*. It seems unlikely that multiple /r/s with two full syllables intervening pose much difficulty for repetition, either on an

articulatory or neural level. We also note that the serialization and articulatory effort theories do not seem to predict insertion of extra /r/s, which occurs both in real words (*farmiliar*, *phertographer*) and in our experiment.

It is important to recognize that the different explanations are not necessarily contradictory. Multiple functional pressures may conspire to make sound repetition both undesirable to create, and unstable where it does exist. Different pressures may also work together: for example, perceptual and/or speech errors may create a pool of variant forms within the speech community, after which grammatical constraints grounded in processing considerations guide the selection of preferred variants among the possibilities.

However, merely acknowledging multiple possible causes does not bring us much closer to solving the ‘actuation’ problem: understanding why dissimilation arises in a particular language at a particular time, and why it takes very different forms in different languages. Nor is it clear whether all the ways that speakers avoid sound repetition (i.e., through diachronic change/dropping of segments, synchronic regular alternations, or statistical underrepresentation in the lexicon) have the same cause, or whether superficially similar repetition-avoidance phenomena may arise and be enforced in different ways.

Of the theories, we argue that Ohala’s perceptual hypercorrection approach holds the most promise for explaining the language-specific character of dissimilation (e.g., why /r/ dissimilates to /l/ in Romance but deletes in English). This theory proposes a very direct and testable link between language-specific phonetics and language-specific diachronic changes.

4. Conclusion. We have shown that in a laboratory task, listeners who are asked to write nonce word are more likely to omit an /r/ in the spelling of a word that contains an additional /r/. We argue that this is a perceptual error, caused by phonetic ambiguity as to how many /r/s are present. The patterns of /r/-dropping produced in this task resemble real-life American r-dissimilation, as seen in words like *surprise*. We believe that perceptual hypercorrection is one of the causes of r-dissimilation, and that perceptual errors play a crucial role in actuation of the change.

References

- Abrego-Collier, Carissa. 2013. Liquid dissimilation as listener hypocorrection. *Proceedings of the Berkeley Linguistics Society (BLS)* 37. 3–17. <https://doi.org/10.3765/bls.v37i1.3195>.
- Carnoy, Albert J. 1918. The real nature of dissimilation. *Transactions and Proceedings of the American Philological Association* 49. 101–113.
- Frisch, Stefan A. 2004. Language processing and segmental OCP effects. In Bruce Hayes, Robert Kirchner & Donca Steriade (eds.), *Phonetically based phonology*. 346–371. Cambridge: Cambridge University Press.
- Gallagher, Gillian. 2010. Perceptual distinctness and long-distance laryngeal restrictions. *Phonology* 27(3). 435–480. <https://www.jstor.org/stable/i40044281>.
- Garrett, Andrew & Keith Johnson. 2013. Phonetic bias in sound change. In Alan C. L. Yu (ed.), *Origins of sound change: Approaches to phonologization*. 51–97. Oxford: Oxford University Press.
- Hall, Nancy. 2009. R-dissimilation in American English. Manuscript, California State University Long Beach.
- Harrington, Jonathan, Felicitas Kleber & Mary Stevens. 2016. The relationship between the (mis)-parsing of coarticulation in perception and sound change: Evidence from dissimilation and language acquisition. In Anna Esposito, Marcos Faundez-Zanuy, Antonietta M.

- Esposito, Gennaro Cordasco, Thomas Drugman, Jordi Solé-Casals & Francesco Carlo Morabito (eds.), *Recent advances in nonlinear speech processing*. 15–34. Dordrecht: Springer.
- Hempl, George 1893. Loss of r in English through dissimilation. *Dialect Notes* 1. 279–81.
- Heinrich, Antje, Yvonne Flory & Sarah Hawkins. 2010. Influence of English r-resonances on intelligibility of speech in noise for native English and German listeners. *Speech Communication* 52(11–12). 1038–1055. <https://doi.org/10.1016/j.specom.2010.09.009>.
- Heid, Sebastian & Sarah Hawkins. 2000. An acoustical study of long-domain /r/ and /l/ coarticulation. *Proceedings of the 5th seminar on speech production: Models and data*. 77–80.
- Martin, Andrew T. 2007. *The evolving lexicon*. Los Angeles, CA: University of California dissertation.
- Meringer, Rudolf & Karl Mayer. 1895. *Versprechen und Verlesen: Eine Psychologisch-Linguistische Studie*. Stuttgart: Göschensche Verlagsbuchhandlung.
- Meringer, Rudolf. 1908. *Aus dem Leben der Sprache: Versprechen, Kindersprache, Nachahmungstrieb*. Berlin: Behr.
- Mondorf, Britta. 2003. Support for more-support. *Topics in English Linguistics* 43. 251–304.
- Ohala, John. J. 1981. The listener as a source of sound change. In Carrie S. Masek, Roberta A. Hendrick & Mary Frances Miller (eds.), *Papers from the parasession on language and behavior*. 178–203. Chicago: Chicago Linguistic Society.
- Ohala, John. J. 1993. The phonetics of sound change. In Charles Jones (ed.), *Historical linguistics: Problems and perspectives*. 237–278. London: Longman.
- Hermann Osthoff & Karl Brugman. 1878. Preface to Morphological Investigations in the Sphere of the Indo-European Languages I. In W. P. Lehmann (ed.), *A reader in nineteenth century historical Indo-European linguistics*. 197–209. Bloomington: Indiana University Press.
- Ozburn, Avery. 2016. Investigating the perceptual hypocorrection hypothesis with sibilant harmony. Paper presented at Conference on Laboratory Phonology 15, July 2016, Cornell, New York.
- Paul, Hermann. 1880. *Principien der Sprachgeschichte*. Halle: Max Niemeyer.
- Posner, Rebecca. 1961. *Consonantal dissimilation in the Romance languages*. Oxford: Blackwell.
- Stephenson, Edward A. 1956. Linguistic resources of the Southern Historical Collection. *American Speech* 31. 271–277.
- Tunley, Alison. 1999. *Coarticulatory influences of liquids on vowels in English*. Cambridge, UK: University of Cambridge dissertation.
- Walter, Mary Ann. 2007. *Repetition avoidance in human language*. Cambridge, MA: Massachusetts Institute of Technology dissertation.
- West, Paula. 1999a. The extent of coarticulation in English liquids: An acoustic and articulatory study. *Proceedings of International Congress of Phonetic Sciences (ICPhS)* 3. 1901–1904.
- West, Paula. 1999b. Perception of distributed coarticulatory properties of English /l/ and /r/. *Journal of Phonetics* 27(4). 405–426. <https://doi.org/10.1006/jpho.1999.0102>.