Analogy in the emergence of intrusive-*r* in English

MÁRTON SÓSKUTHY
University of Edinburgh

ABSTRACT

This paper presents a novel approach to the phenomenon of intrusive-*r* in English based on analogy. The main claim of the paper is that intrusive-*r* in non-rhotic accents of English is the result of the analogical extension of the *r*∼zero alternation shown by words such as *far*, *more* and *dear*. While this idea has been around for a long time, this is the first paper that explores this type of analysis in detail. More specifically, I provide an overview of the developments that led to the emergence of intrusive-*r* and show that they are fully compatible with an analogical approach. This includes the analysis of frequency data taken from an 18th century corpus of English compiled specifically for the purposes of this paper and the discussion of a related development, namely intrusive-*l*. To sharpen the predictions of the analogical approach, I also provide a mathematically explicit definition of analogy and run a computer simulation of the emergence of the phenomenon based on a one million word extract from the 18th century corpus mentioned above. The results of the simulation confirm the predictions of the analogical approach. A further advantage of the analysis presented here is that it can account for the historical development and synchronic variability of intrusive-*r* in a unified framework.

1 INTRODUCTION

The phenomenon of intrusive-*r* in various accents of English has inspired a large number of generative analyses and is surrounded by considerable controversy, mainly because of the theoretical challenges that it poses to Optimality Theory and markedness-based approaches to phonology (e.g. McCarthy 1993; Harris 1994; Halle & Idsardi 1997; Baković 1999; Uffmann 2007). I believe that this line of research is crucially misguided: its narrow focus on minor details of analysis leads to a general lack of interest in the complex interactions that make this phenomenon particularly intriguing. Even more problematically, many of these analyses make dubious claims about the empirical aspects of intrusive-*r* and disregard the results of existing quantitative studies. Two areas that have been particularly neglected in generative discussions of the phenomenon are its historical development and its extreme variability. The general indifference with respect to these areas stems directly from the underlying principles of the generative programme, according to which the primary goal of linguists is to construct *synchronic* models of *competence*, which therefore do not have to deal with issues of diachrony or performance. However, these restrictions have not proven particularly felicitous in the case of intrusive-*r*, whose apparent unnaturalness has led many researchers to claim that it is synchronically arbitrary (McCarthy 1991, 1993; Blevins 1997; Halle & Idsardi 1997; McMahon 2000), thereby implicitly acknowledging diachrony as a potential source of explanation for its behaviour in present-day accents. In this paper, I show that this 'diachrony-as-a-last-resort' approach is insufficient: the facts related to intrusive-*r* can only be fully understood by taking an explicitly diachronic approach and exploring the history of the phenomenon in detail.

The main claim of this paper is that the pattern of intrusion seen in Southern British English and other intruding accents is the result of a process of analogical extension. Several different types of argument are presented in support of this hypothesis. First of all, I show that the analogical approach makes correct predictions about the diachronic development of the phenomenon. This includes a detailed overview of the frequency distribution of word classes related to intrusive-*r*, which are investigated in a corpus of 18th century English compiled specifically for the purpose of this paper. The frequency arguments receive further support from a related development in a number of American accents, namely intrusive-*l*. I also demonstrate that once a formally explicit definition of analogy is adopted, the facts about the development of intrusive-*r* and its variability can be accounted for in a unified way. This is corroborated by a computer simulation of the emergence of the phenomenon, which takes a portion of the 18th century corpus mentioned above as its input, and eventually produces an accent with a variable process of intrusion.

It will be useful to clarify the use of a number of key terms in this paper. Words which contain an inetymological *r* in intruding accents of English are referred to as *r*-LESS (because of their lack of an *r* before the appearance of intrusive-*r*), and words with an etymological *r* as *r*-FUL. Historical accents which had developed the conditions necessary for the emergence of intrusive-*r* are termed PRE-INTRUSION ACCENTS. Preconsonantal and prepausal instances of *r* are simply referred to as CODA-*r*, although I would like to emphasise that I do not make any assumptions about the actual syllabic status of this consonant in English and use this term purely for convenience.[1]

The structure of the paper is as follows. In Section 2, I provide a preliminary outline of the analogical approach to intrusive-*r* and present its main predictions. Section 3 sets these predictions against the empirical aspects of intrusive-*r* and finds that while the particular conditions in pre-intrusion dialects are fully compatible with an analogical account, the variability of the phenomenon cannot easily be explained without elaborating on the notion of analogy itself. This task is taken up Section 4, which presents an overview of analogical models and arrives at a more explicit formulation of analogy based on tokens of use. Finally, Section 5 develops the token-based approach into a computer simulation and shows that it makes accurate predictions about the historical development and the variability of intrusive-*r*. I conclude the paper with a brief summary of its main points.

2 PRELIMINARY ANALYSIS

The term intrusive-*r* refers to an *r*∼zero alternation at the end of *r*-less words; the variant with a final *r* appears before a vowel and the one without *r* before a consonant or a pause. According to most traditional accounts, intrusive-*r* only involves words with a final [ɑː], [ɔː] or schwa (e.g. *spa*, *law* and *pizza*, respectively). The epenthetic consonant may occur across words (e.g. *spa*[r] *is*, *law*[r] *and order*, *idea*[r] *of* and word-internally as well (e.g. *withdraw*[r]*al*, *saw*[r]*ing*). It is important to note that intrusive-*r* only appears in non-rhotic dialects, that is, dialects in which *r*-ful words also show a final *r*∼zero alternation (e.g. *scar*, *lore*, *Peter*). The alternation in these words is the result of a historical process of *r* Dropping before consonants and at the end of the word.

The main argument of this paper is that intrusive-*r* appeared in *r*-less words under the analogical influence of *r*-ful words. To put it slightly differently, the alternating pattern of *r*-ful words was analogically extended to the *r*-less group, resulting in a complete merger between the two classes, illustrated in (1) below (the shading illustrates the extent of the merger):

(1)

	R-FUL		R-LESS		R-FUL	R-LESS
_{C, ‖}	V#		V#	⟹	V#	V#
_V	Vr#	→	V#		Vr#	Vr#

This insight also forms the basis of several previous analyses of the phenomenon, among them Jones (1964), Gimson (1980), Gick (1999, 2002), and Bermúdez-Otero (2005). However, while I believe that this approach is intuitively appealing, analogy has little – if any – explanatory power unless one specifies the exact conditions under which a pattern can be extended and demonstrates that these conditions are present in the language where the extension is suggested to occur. In the present case, this means (i) identifying the situations in which extension is likely to occur; (ii) giving a precise description of what qualifies as a potential analogical source in such a situation; and (iii) showing that such a situation arguably held in pre-intrusion dialects of English with the *r*-ful class being a suitable analogical source.

As for (i), most contemporary approaches to analogical extension assign a crucial role to similarity (Skousen 1989; Albright & Hayes 2003; Albright 2009): the likelihood of the extension of a pattern is a function of the similarity between the analogical source and the analogical target; the more similar they are, the more likely it is that the extension will occur. Turning now to (ii), the likelihood of a pattern to serve as the source of the extension is usually claimed to be proportional to its frequency (Bybee 2001), which means that the direction of the extension is determined by the relative frequencies of the two patterns: the analogical source will normally be of higher frequency than the analogical target. This means that the analogical approach makes two crucial predictions about *r*-ful and *r*-less words in pre-intrusion dialects:

Prediction 1 *R-ful and r-less words are similar.*

Prediction 2 *R-ful words are more frequent than r-less words.*

It should be noted that the exact role of similarity and frequency in analogical models is left unspecified for the moment being – I will simply assume that the consensus of the recent analogical literature on the importance of these concepts is sufficient to treat them as essential components of an analogical account. This vagueness is remedied in Sections 4 and 5, where these notions are substantiated and formalised within a computationally explicit framework.

3 EMPIRICAL ASPECTS OF INTRUSIVE-R

3.1 The historical development of intrusive-r

This section provides an overview of the historical developments related to intrusive-*r* and shows that the predictions of the analogical approach are borne out by the data. It is reasonable to assume a similar set of conditions to have held in all of the dialects where intrusive-*r* emerged (at least with respect to intrusion); therefore, I focus my attention on a single dialect, Southern British English (SBE), and suggest that the same points could also be made for other dialects with intrusive-*r*. Since the first evidence of intrusive-*r* in SBE comes from Sheridan's *A Course of Lectures on Elocution* from 1762 (Sheridan 1762/1803), I will assume that the standard dialect spoken in the south of England in the mid-18th century is an example of a pre-intrusion dialect.

Let us first take a look at Prediction 1, which concerns the similarity between *r*-ful and *r*-less words. It is clear that in present-day Southern British English *r*-ful and *r*-less words share an important structural feature: the set of final vowels appearing in preconsonantal and prepausal allomorphs of *r*-less words (i.e. [ə, ɔː, ɑː]) is a subset of the set of final vowels appearing in *r*-ful words in the same environment (i.e. [ə, ɔː, ɑː, ɜː]). The analogical approach to intrusive-*r* requires this structural similarity to be present in pre-intrusion dialects as well, which means that it needs to be shown that SBE acquired this particular distribution of final vowels no later than the middle of the 18th century. There are several pieces of evidence that suggest that this might well have been the case. The single most important factor in the emergence of the partial overlap between the two relevant classes of words is the loss of *r* in preconsonantal and prepausal position, which created the word-final *r*∼zero alternations exhibited by *r*-ful words in present-day English:

(2) *r* Dropping: r → Ø/ __{C, ‖}

__ ‖		__C		__V	
[wɔː‖]	'war'	[wɔːwɒz]	'war was'	[wɔːrɪz]	'war is'
[stɑː‖]	'star'	[stɑːlaɪt]	'starlight'	[stɑːrɒn]	'star on'
[bɛtə‖]	'better'	[bɛtəðæn]	'better than'	[bɛtərɪn]	'better in'

While Wells (1982) dates this change after 1750, Lass (2000) and McMahon (2000) argue that the decline of coda-*r* started much earlier, perhaps in Early Modern English, with the weakening of preconsonantal and prepausal *r*, and was already 'under way, producing variants in the speech community, before 1700' (McMahon 2000: 234). For a detailed overview of the orthoepic evidence the reader is referred to McMahon (2000: 237-241). It is important to note that the historical sources do not point to a complete disappearance of coda-*r* in SBE:

there is a marked lack of agreement among 18th century authors as to whether coda-*r* is pronounced or not, which suggests that *r* Dropping was variable at this stage. However, it is fair to assume that a considerable proportion of coda-*r*'s was now being dropped, creating a sufficient amount of overlap between the *r*-ful and the *r*-less classes to serve as the basis of analogical extension (see Section 5 for a more detailed discussion of what is meant by a 'sufficient amount of overlap').

There are two further changes that played an important role in shaping the distribution of vowels before *r* termed Pre-*r* Lengthening and Pre-*r* Broadening by Wells (1982). These are illustrated in (3) and (4) below:

(3) Pre-*r* Breaking:[2] Ø → ə / {iː, eː, oː, uː, aɪ, aʊ}__ r

[biːr]	>	[bɪər]	'beer'
[tʃeːr]	>	[tʃɛər]	'chair'
[moːr]	>	[mɔːər]	'more'
[ʃuːr]	>	[ʃʊər]	'sure'
[faɪr]	>	[faɪər]	'fire'
[taʊr]	>	[taʊər]	'tower'

(4) Pre-*r* Lengthening: {ɑ, ɔ, ɜ} → {ɑː, ɔː, ɜː} / __ r{C, #}

[bɑrd]	>	[bɑːrd]	'bard'
[hɔrn]	>	[hɔːrn]	'horn'
[bɜrd]	>	[bɜːrd]	'bird'

Since both of these changes were conditioned by the following *r*, it is clear that they had to predate the emergence of complete non-rhoticity. This argument is also supported by the historical record: Pre-*r* Breaking seems to have been a long and gradual process, starting already in the 16th century (see Jespersen 1909; Jones 1989), and Pre-*r* Lengthening was also underway from at least the beginning of the 18th century (see McMahon 2000: 235-236). It should be noted that while both of these changes increased the amount of overlap between the *r*-ful and the *r*-less classes, only Pre-*r* Lengthening was crucial to the creation of the overlap itself. That is to say, the changes in vowel quantity (and possibly also vowel quality; cf. MacMahon 1998) occasioned by Pre-*r* Lengthening were a prerequisite for the merger of *r*-ful words which ended in [ɔr]/[ɑr] before the 18th century (e.g. *war* EMoE [wɔr], *star* EMoE [stɑr]) and *r*-less words in [ɔː]/[ɑː] (e.g. *law* [lɔː], *Ma* [mɑː]). On the other hand, the loss of coda-*r* would have created a large number of schwa-final forms even if Pre-*r* Breaking had not taken place (e.g. *better* [bɛtə], *author* [ɔːθə], *altar* [ɔːltə]; see below for more detail).

It is remarkable that all the accents where intrusive-*r* has emerged share these features with SBE: all intruding accents are non-rhotic, they all show the effects of Pre-*r* Lengthening, and they all have centring diphthongs in *r*-ful words (although

many of these accents have subsequently monophthongised [ɛə], [ɔə] and [ʊə]). This can be interpreted as further evidence for the analogical approach: intrusive-*r* only emerges in accents where there is a phonetic overlap between the *r*-ful and the *r*-less classes (i.e. where they have identical final vowels). Even more interestingly, the number of non-rhotic accents without intrusion is conspicuously low. Southern American English and South African English appear to be the only accents where non-rhoticity does not entail intrusion, if we disregard speakers of various other non-rhotic accents who consciously avoid intrusion. Incidentally, these accents also share another important feature, namely that etymologically *r*-ful words are more or less consistently realised without a final *r* even in prevocalic position (Wells 1982; McDavid 1958). Once again, these observations receive a straightforward interpretation if we take analogy to be the source of intrusive-*r*: in these accents, *r*-ful words have a non-alternating pattern, which cannot yield an *r*-zero pattern in *r*-less words through analogical extension.

Let us now turn to Prediction 2, which is about the frequency distribution of *r*-ful and *r*-less words. I have suggested that analogical extension only occurs if the source of the pattern is of higher frequency than the target of the extension. To test whether this relationship held between the *r*-ful and the *r*-less classes in pre-intrusion SBE, I compiled a 2 million word phonetically annotated corpus of early and mid 18th century English (henceforth CE18). The corpus consists of several 18th century novels (among them Samuel Richardson's *Clarissa* and Daniel Defoe's *Robinson Crusoe*) and all issues of *The Spectator* between 1711 and 1714. Foreign sentences and medium-specific features such as chapter headings and signatures were removed from the text. The automatic phonetic annotation of the text was based on the transcriptions of the CELEX database (Baayen et al. 1995); a number of transcriptions were added manually. While I believe that this corpus can provide us with more accurate details about the frequency distributions of the relevant word classes in 18th century English than any present-day corpus, I acknowledge that it has a number of drawbacks. For instance, all the sources reflect the language use of the educated upper class, rather than that of the lower class, although intrusive-*r* quite possibly originated in the language of the latter (cf. Sóskuthy 2009). Nevertheless, the variety of English that dominates the corpus is probably the closest that we can get to the varieties in which intrusive-*r* first appeared. A further problem with CE18 is that the transcriptions – being based on CELEX – reflect present-day pronunciations rather than 18th century ones. Once again, this may not be such a serious disadvantage, given that the phonological differences between 18th century English and Present-day English are relatively minor and do not involve the main characteristics of the lexical classes that this analysis is based on.

The token frequencies of *r*-ful and *r*-less words are presented in Table 1 – the reason for choosing token frequencies over type frequencies will become clear

	R-LESS	R-FUL	RATIO
ə#	1,553	99,979	1:64.38
ɔː#	1,487	51,871	1:34.88
ɑː#	112	9,397	1:83.90
SUM	3,152	161,149	1:51.13

Table 1: *The token frequencies of* r-*ful and* r-*less words in the CE18 corpus*

from the discussion in Section 4.2. The size of the *r*-ful class is two orders of magnitude greater than that of the *r*-less class, which confirms Prediction 2: the proposed analogical source is of considerably higher frequency than the analogical target. It should be noted that words in centring diphthongs make a smaller contribution to the *r*-ful class than words ending in a lone [ə]: there are 66,721 *r*-ful tokens which end in a lone [ə] and only 33,258 tokens which end in a centring diphthong (including [aɪə] and [aʊə]). This corroborates the assertion made above that Pre-*r* Breaking was not essential in the emergence of the overlap between the two relevant word classes – *r*-ful words in a lone schwa and *r*-less words in schwa would have identical endings in preconsonantal and prepausal position even if Pre-*r* Breaking had not taken place.

One final piece of evidence in support of the analogical approach comes from a related but independent development in certain Mid-Atlantic dialects of American English, namely intrusive-*l* (see Gick 1999, 2002). Intrusive-*l* shows a very similar distribution and development to intrusive-*r*: it appears in words with a final [ɔː] (but not [ə] and [ɑː]) when the following word is vowel-initial (e.g. *draw it* [drɔːlɪt] vs. *draw them* [drɔːðəm]), and is only found in dialects where *l* has been lost in preconsonantal and prepausal position. This suggests that intrusive-*l* might also be a case of analogical extension based on the partial merger of previously *l*-ful and *l*-less forms (e.g. *drawl* and *draw*). If this is the case, we expect to find the same asymmetric frequency distribution for *l*-ful and *l*-less forms as for *r*-ful and *r*-less forms. This prediction is partially borne out by the data, as can be seen in Table 2 (the frequency counts are taken from the CELEX corpus, as intrusive-*l* seems to have emerged significantly later than intrusive-*r*). What is particularly striking here is that the frequency distribution necessary for the extension of the *l*-ful pattern (i.e. the analogical source is of higher-frequency than the analogical target) only holds for words in [ɔː], which may well be the reason why words in [ə] and [ɑː] have not developed intrusive-*l*. Moreover, the frequency differences between *l*-ful and *l*-less words are markedly smaller than those between *r*-ful and *r*-less words. This may explain why intrusive-*l* is less systematic and wide-spread than intrusive-*r* (cf. Gick 2002): a smaller difference in frequency entails a slower

	L-LESS	L-FUL	RATIO
ə#	45,932	41,282	1:0.90
ɔː#	12,219	110,874	1:9.07
ɑː#	4,237	704	1:0.16

Table 2: *The token frequencies of* l*-ful and* l*-less words in the CELEX corpus*

process of extension. While this evidence is clearly circumstantial with respect to intrusive-*r*, the fact that analogy provides a unified explanation for two independent processes of intrusion in English and makes valid predictions for both is a strong argument for adopting an analogical approach.

To conclude this section, let us sum up its main points. It has been shown that intrusive-*r* conforms to the two main predictions of the analysis presented in Section 2: the *r*-ful class and the *r*-less class share essentially the same set of final vowels and the *r*-ful class has considerably higher token frequency than the *r*-less class. Moreover, we have also seen that the analogical approach can provide a straightforward explanation for a number of related issues: the absence of intrusive-*r* in Southern American English and South African English and the development of intrusive-*l* in Mid-Atlantic varieties of English.

3.2 The variability of intrusive-r

In the preceding section, a number of simplifying assumptions were made to allow for a more streamlined presentation of the issues relevant to the development of intrusive-*r*. More specifically, I presented the relevant word classes and the phenomenon of intrusive-*r* itself as if they behaved categorically and referred to variation only occasionally. However, it appears that the actual situation is somewhat more complicated: *r* Dropping is likely to have been a gradual process creating a considerable amount of variation in *r*-ful words (cf. above) and intrusive-*r* is variable even in present-day accents, which suggests that it did not appear overnight in pre-intrusion accents. Therefore, in this section I discuss the issue of variability and point out its relevance to various analyses of intrusive-*r*.

First of all, there is some evidence to suggest that the accents in which intrusive-*r* emerged were only partially non-rhotic. Hay & Sudbury (2005) is a rather unusual study which examines the incidence of linking and intrusive-*r* in the speech of several generations of New Zealanders born between 1850 and 1930, based on a collection of audio recordings, most of which were made around 1940. They find a high degree of variability for both *r* Dropping and intrusive-*r* in their sample, which suggests that even a pre-intrusion accent can be partially rhotic. Moreover, their study shows that the incidence of intrusive-*r* is significantly correlated with

the speakers' degree of rhoticity: '[i]ntrusive /r/ increases as rhoticity declines' (Hay & Sudbury 2005: 813). While this seems to be a serious challenge to the analogical approach, which assumes that intrusive-*r* is a direct consequence of the partial merger between the *r*-ful and the *r*-less occasioned by the loss of coda-*r*, it will be shown in Section 5 that a certain amount of variation among *r*-ful words is easily handled by an token-based analogical framework.

Moreover, empirical studies of intrusive-*r* in present-day English consistently find a certain amount of variation in the realisation of intrusive-*r* (Bauer 1984; Foulkes 1998; Sóskuthy 2009; Hay & MacLagan to appear). Some of these studies also point out that this variation is governed by grammatical and lexical factors such as morphological boundary strength (i.e. more productive suffixes such as *#ing* and *#ish* attract higher rates of intrusion; Hay & MacLagan to appear), the identity of the final vowel (Bauer 1984; Hay & MacLagan to appear) and the identity of the target word (Sóskuthy 2009; Hay & MacLagan to appear). These findings – together with Hay & Sudbury's (2005) results – also suggest that the emergence of intrusive-*r* was gradual: it would be unusual to find that such a highly variable pattern emerged without a relatively long period of transition.

I believe that the results presented above speak strongly against generative treatments of intrusive-*r*. First of all, most generative accounts assume that this phenomenon is more or less categorical: for instance, Blevins (1997) claims that intrusive-*r* is an 'exceptionless [process] which [applies] both to native vocabulary items and loans' (ibid. 247). While it is true that generative models are forced to treat the phenomenon as categorical due to their inability to capture meaningful patterns of variation, all empirical studies report a great deal of variability in the production of intrusive-*r*. What is even more problematic for generative models is that this variation is to be controlled by grammatical and lexical factors. There is no way for theories adopting a strictly categorical view of grammar to account for these sophisticated patterns of variation without radically departing from the principles that underlie the generative programme.

The question arises as to whether the above observations about variation can be accounted for within an analogical framework. The simple analogical scenario sketched in the preceding section does not make any particular predictions about variation in the production of individual words – in fact, it does not seem to fare much better with respect to variation than generative models do. However, once we make the notion of analogical extension more explicit, it becomes possible to account for at least some of the observed patterns of variation. This task will be taken up in the next section.

4 WHAT IS ANALOGY?

In the preceding section, I reviewed two different sets of evidence and concluded that (i) intrusive-*r* in SBE is likely to be the result of analogical extension and (ii) it shows a great deal of variation. The main goal of this section is to suggest some ways in which these two different observations can be integrated in a single analogical framework. To achieve this goal, it will be necessary to give a more explicit definition of analogy itself and briefly review previous approaches to analogical extension (4.1). It will be shown that the structure of existing analogical frameworks makes them incapable of handling variation at the level of individual words, and a different approach will be suggested (4.2).

4.1 Previous approaches to analogy

While the term analogy is used in a variety of ways in the literature (see Hock 2003 for an overview), this paper focusses on one particular mechanism, which seems to serve as the basis of most computationally implemented models of analogy, namely FOUR-PART ANALOGY. Four-part analogy consists in the extension of a certain relationship between a pair of forms to another pair of forms, where the members of the two pairs bear the same structural or semantic relationship to each other. An example is given in (5) below:

(5)

	[singular]		[plural]	
BOW	[baʊ]	∼	[baʊz]	
		↓		
COW	[kaʊ]	∼	?	(= [kaʊz] < [kaɪn])

The four edges of the analogical rectangle will be referred to as follows: the SOURCE (BOW), the TARGET (COW), the KNOWN ENVIRONMENT ([singular]) and the GIVEN ENVIRONMENT ([plural]). The corners of the rectangles can be identified by referring to the two edges that meet there: for instance, [baʊ] is the source in the known environment and [kaʊz] (the form that we obtain through analogical extension) is the target in the given environment (this will also be referred to as the GIVEN FORM). The particular relationship that is extended in (5) can be described as $\{x \sim x + [z]\}$. This relationship clearly yields [kaʊz] when applied to the target in the known environment, that is, [kaʊ].

This type of analogy can also be used to model the extension of the *r*-ful pattern to an *r*-less word:

(6)

	_C		_V	
DEAR	[dɪə]	∼	[dɪər]	
		↓		
IDEA	[ajdɪə]	∼	?	(= [ajdɪər] < [ajdɪə])

As it is pointed out by Albright (2009), this type of formalism does not impose any restrictions on the choice of the analogical source: in the example in (6), the lexeme DEAR is used, but other lexemes, such as MARIA, CAT or SMURF could equally well have been used, in which case no change would have taken place (as these lexemes do not show an alternating *r*∼zero pattern). This is clearly problematic: the transition from the analogical target to the analogical source is arguably guided by frequency and similarity, as it has been noted in Section 3.1. Moreover, it is not clear what the exact role of four-part analogy is within language change: it seems to be a useful descriptive device, but should we assume such analogical extensions to take place in individual speakers as well? And if yes, do these extensions occur during speech production or speech perception? The four-part notation cannot answer these questions in itself.

Most computationally implemented models of analogy take a somewhat simplified version of the four-part analogical mechanism as their starting point and use a number of extra mechanisms to ensure that both similarity and frequency have an effect on the choice of the analogical source. It will be useful to take a brief look at a particular class of such models, namely INSTANCE-BASED LEARNERS, some examples of which are the GENERALIZED CONTEXT MODEL (GCM; Nosofsky 1986, 1988), ANALOGICAL MODELING (AM; Skousen 1989; Skousen et al. 2002) and the TILBURG MEMORY-BASED LEARNER (TiMBL; Daelemans et al. 2007).[3]

Instance-based learners are based on the assumption that the behaviour of a given item can be determined by comparing it to similar items within the dataset. The dataset for an instance-based learner could consist of a list of phonetically transcribed types from the lexicon of English, where each type is associated with a particular behaviour in prevocalic position, as exemplified in Table 3. The types are represented as a set of variables, which, in this case, are the last five sounds of each occurrence ('=' means non-specification for a given feature). Instance-based learners can use this dataset to predict the behaviour of any item that is specified using the same variables. This could be a new item, which is not present in the original dataset (this would be similar to a learner trying to establish a certain pattern for a nonce-form) or an item from the dataset itself (as in the case of analogical extension, where an existing pattern is replaced by a new one).

The model's prediction is based on the behaviour of items that are similar to the given form – TiMBL achieves this by building a restricted analogical set,

LEXEME	VARIABLES	BEHAVIOUR
bread	=, b, ɹ, ɛ, d	{+ø}
spin	=, s, p, ɪ, n	{+ø}
city	=, s, ɪ, t, i	{+ø}
idea	a, ɪ, d, ɪ, ə	{+ø}
law	=, =, =, l, ɔː	{+ø}
four	=, =, =, f, ɔː	{+r}
better	=, b, ɛ, t, ə	{+r}
star	=, =, s, t, ɑː	{+r}

Table 3: *Dataset for selection of patterns of alternation in English*

which consists of the nearest neighbours of the given form, and the GCM by giving a greater weight to items that are more similar to the given form in determining the probability of an item to serve as the analogical source (AM also works with a restricted analogical set, but the membership within this set is only indirectly influenced by similarity). The precise calculation of similarity values differs from model to model, but in most cases it is a function of the number of overlapping variables, where certain variables can have a greater influence than others. Thus, IDOL [a, ɪ, d, ə, l] and DEAR [=, =, d, ɪ, ə] both share three variables with IDEA [a, ɪ, d, ɪ, ə], but the last variable can be given a greater weight in determining similarity values, as it is more relevant to the task at hand than, say, the first variable.[4]

Frequency influences the predictions of instance-based learners in a less direct way. The likelihood of any individual form to serve as the analogical source or be included in the analogical set is solely determined by its similarity to the given form. However, since a high-frequency behavioural pattern is necessarily better represented in the dataset than a low-frequency one, it will have a greater chance of influencing the outcome of the prediction, provided that the items are relatively evenly distributed in the feature space defined by the variables. For instance, if there are 90 items with behaviour *A* and only 10 items with behaviour *B*, any random point in the feature space will be likely to be surrounded by a majority of items with behaviour *A*. The only scenario in which behaviour *B* can have any significant effect on the outcome of the prediction is when the items with behaviour *B* form a tight group (sometimes referred to as a 'gang'; cf. Bybee 2001) within the feature space, that is, when they are consistently more similar to each other than to items with behaviour *B*. Figure 1 provides an illustration of a dataset where the distribution of the items in the feature space is independent of their behaviour (left) and another dataset where items with a low-frequency pattern form a tight group (right).

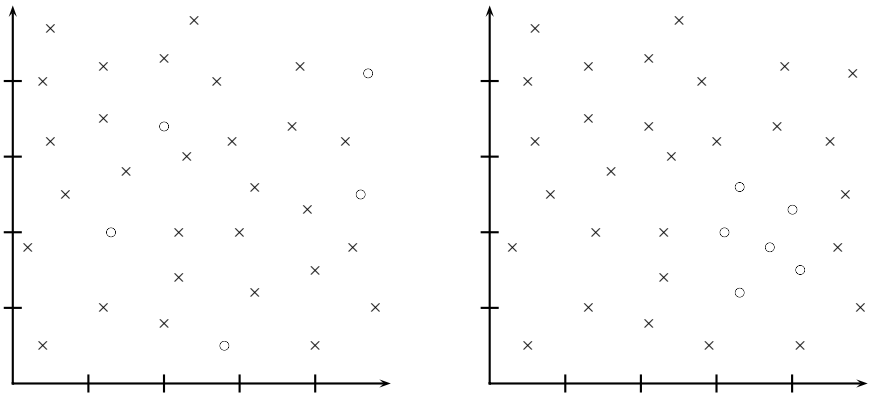

Figure 1: *Left panel: items with the two behavioural patterns are evenly distributed in the feature space; Right panel: items with the low-frequency pattern form a tight group.*

This might also help us understand how an instance-based learner could be used to model the emergence of intrusive-*r*. Pre-intrusion accents of English exemplify the evenly distributed scenario (i.e. *r*-less words are randomly dispersed among members of the *r*-ful class), whereas accents in which the partial merger between the *r*-ful and the *r*-less classes did not take place exemplify the second scenario, with *r*-less words forming a tight group. Therefore, in a pre-intrusion accent the outcome of the prediction will be more strongly influenced by the *r*-ful pattern than by the *r*-less one owing to the higher frequency of the former; *r*-ful forms will simply have a greater chance of being included in the analogical set or being chosen as the analogical source. The model will tend to predict an *r*-ful pattern of behaviour even for *r*-less words, that is, analogical extension will take place. This is confirmed by TiMBL, which extends the *r*∼zero pattern to the majority of *r*-less words (75% to 100% depending on the parameter settings) when confronted with a dataset based on a pre-intrusion dialect such as the one in Table 3. However, in accents where the *r*-ful and the *r*-less classes are fully distinguishable, words within the *r*-less class will be more similar to each other than to words within the *r*-ful class. This similarity will counterbalance the higher frequency of the *r*-ful pattern and result in the retainment of the distinction between the two classes.

While instance-based learners can clearly capture some crucial aspects of the analogical extensions that led to the emergence of intrusive-*r*, it has to be pointed out that their success hinges on a considerable simplification: they do not make a distinction between the known environment and the given environment and they assign a single pattern of behaviour to each item. By doing so, they essentially reduce the problem of analogical extension to a simple categorisation task: a stimulus represented by a feature vector has to be assigned a category label, which is a certain pattern of behaviour in this case. In the following section, I show that this

LEXEME	VARIABLES	BEHAVIOUR
better	=, b, ɛ, t, ə	{+r}
better	=, b, ɛ, t, ə	{+ø}
better	b, ɛ, t, ə, r	{–r}
better	b, ɛ, t, ə, r	{+ø}

Table 4: *Dataset for selection of patterns of alternation in English*

approach is problematic in several respects.

4.2 Types and tokens

The simplification that serves as the basis of instance-based learners comes at a price: we are forced to discard all information about variation below the word level. Each type is assigned a single feature vector and a single pattern of behaviour, which is either *r*-ful or *r*-less in this particular case. As a result, a number of arbitrary decisions have to be made, which lead to considerable conceptual and empirical difficulties.

First of all, as types are abstractions over a set of tokens, they often cannot be associated with a unique representation. Choosing the citation forms of the types in Table 3 was a completely arbitrary decision; I might as well have used their prevocalic forms, in which case there could be no analogical extension (as *r*-ful and *r*-less words are distinct in prevocalic position in pre-intrusion accents). In fact, it might be just simply impossible to assign any phonetic representation to types which have several alternants. If a type is a collection of properties shared by a number of tokens, the phonetic forms of the individual tokens are arguably not part of it when they differ from token to token.

Another related problem is that types often cannot be associated with a unique behaviour in a natural linguistic setting, having variable outcomes instead. Hay & Sudbury (2005) suggest that *r*-ful words are likely to have been realised variably in pre-intrusion dialects, with *r*-ful and *r*-less productions occurring both in preconsonantal/prepausal and prevocalic position. The only way to include this type of information in an instance-based learner is to have several items represent a single type, each with a different behaviour or feature representation. For instance, an *r*-ful word with two variants in both preconsonantal/prepausal and prevocalic position could be represented as in Table 4. Besides being intuitively unappealing, this method leads to a proliferation of behavioural patterns (cf. the new pattern {–r} in Table 4) and it unnecessarily strengthens the representation of variable types in the dataset as opposed to invariable ones (which can still be represented by a single item). If, on the other hand, one decides to ignore the information about

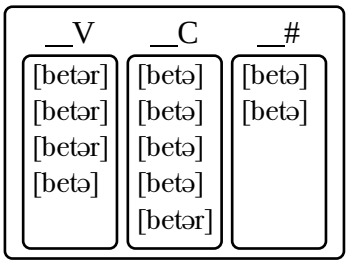

Figure 2: *Exemplar cloud of R-ful word*

variation, an arbitrary choice has to be made between the different patterns when creating the dataset, which leads to the same conceptual problems discussed in the previous paragraph.

What is more, instance-based learners are faced with the same question as the four-part model described in the previous section: it is not clear where the extensions should be assumed to take place. Since the analogical task described in the previous section consisted in the model predicting a production pattern for words on the basis of similar words, it could be argued that the locus of the extensions is the speech production of individual speakers. However, this assumption is somewhat problematic in view of the fact that instance-based learners immediately reclassify almost all *r*-less words as *r*-ful if the input dataset is non-rhotic – this leaves no room for the gradual development implied by the empirical findings in Section 3.2. Thus, while instance-based learners are useful in showing that some type of analogical extension can produce a pattern of intrusion, they are not particularly revealing as to how exactly these extensions occur.

I believe that a more accurate model of the emergence of intrusive-*r* can be constructed by taking an exemplar theoretic model of the lexicon such as the one presented in Bybee (2001) and Pierrehumbert (2001) and combining it with the idea that sporadic changes can accrue over several generations and lead to more robust patterns (Wedel 2004, 2007; Oudeyer 2006). The basic idea in exemplar theory is that linguistic categories such as words and sounds are represented directly by detailed memory traces in an associative network. In this particular case, this means that all tokens of use are stored in the lexicon linked to the specific context (semantic, phonological, social, etc.) in which they are used. Therefore, a model of this type will directly represent variation in so-called clouds of exemplars, as in Figure 2.

Let us assume that this associative network gives rise to a small number of analogical extensions during speech production – in this case, extensions of the *r*-ful pattern. The forms created through such extensions will be passed on to future generations of speakers, who will produce further extensions, thus steadily increasing the proportion of analogically re-modelled forms. All things being

equal, this will eventually result in a large proportion of forms with intrusive-*r* (this is called POSITIVE FEEDBACK in Wedel 2007). Importantly, the speakers' lexicon will contain a great deal of variation at any given point in this process and the emergence of intrusive-*r* will be gradual. In the next section, I provide a more detailed exposition of this model.

5 TOKEN-BASED ANALOGY AND INTRUSIVE-R

In the previous section, several formally more or less explicit approaches to analogy were discussed. It was shown that while they are capable of producing the type of analogical extensions necessary for the modelling of the emergence of intrusive-*r*, they cannot capture the patterns of variation observed in Section 3.2 due to certain restrictions in their structure. At the end of the section I proposed that a model in which the basic units of storage are tokens rather than types could produce more realistic results, but I have not elaborated on the details of this model. Therefore, in this section, I present a formal model of analogy built on tokens, which will serve as the basis of a computer simulation of the emergence of intrusive-*r*.

In the brief description of the token-based approach of analogy above, several assumptions were made which need to be explored in more detail. First of all, I suggested that an associative network of tokens can somehow automatically produce analogical extensions. This is, of course, not the case. It is true that a connectionist network along the lines of Rumelhart & McClelland (1986) is capable of modelling certain types of analogical extension, but this requires the network to be set up in a particular way – a random collection of tokens will not start behaving as a coherent system just because they are connected in a network. In this paper, I use an analogical mechanism superimposed on the network of tokens to produce the required extensions. This mechanism is exposed in more detail below. Moreover, I also claimed that these extensions are sporadic and can accumulate over several generations. This is modelled by using the analogical mechanism to predict productions for the items in the dataset and then taking the predicted productions as the input dataset for the next generation. By repeating this procedure several times, we can see whether analogically re-modelled forms can accumulate in the lexicon to form a more robust pattern (this is called ITERATED LEARNING; cf. Brighton 2003; Kirby et al. 2007).

To make this more explicit, here is an outline of the structure of the simulation. The initial input of the model is a list of tokens from the CE18 corpus, represented as an ordered triplet consisting of the phonetic form of the item, the lexeme the item belongs to and the phonetic environment it appears in (which can be __C, __V or __#, depending on the first sound of the following word). Thus, a precon-

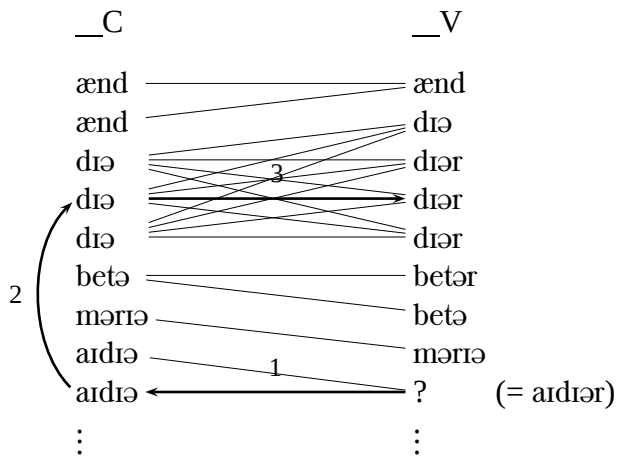

Figure 3: *Constructing a plausible output form for* IDEA *in* __V*; (1): choosing a token of the target in the known environment; (2): choosing an analogical source in the known environment; (3): choosing a token of the analogical source in the given environment.*

sonantal production of IDEA will be encoded as follows: {[ajdɪə], IDEA, __C}. The model goes through each item in the list and predicts a pronunciation for it using the analogical mechanism described below (the predicted pronunciation may or may not be identical to the stored one).[5] These pronunciations are stored in the lexicon of the next generation, which starts its own round once the first generation has produced all the items in the dataset. This process can be repeated indefinitely, but we will see that 30 rounds are sufficient for our purposes.

The crucial step in this process is, of course, the prediction of pronunciations for the items in the dataset. These predictions are based on a four-part analogical mechanism, as shown in Figure 3. Here is a step-by-step description of this process. The input of the analogical prediction is an ordered pair consisting of the lexeme the item belongs to and its environment – in Figure 3, this is {IDEA, __V}. This determines the analogical target (IDEA) and the given environment (__V). To complete the analogical rectangle, we first have to find another environment (the known environment) with at least one token of the target. In our example, the known environment is __C. Now, a random token of the target lexeme is chosen in the known environment (step 1 in Figure 3), which will serve as the basis of our choice of the analogical source in the known environment (step 2). The transition from the analogical target to the analogical source is determined by three factors: (i) similarity to the target in the known environment, (ii) token frequency and (iii) the availability of at least one form belonging to the same lexeme in the given environment. The second factor is crucial, as the next step consists in randomly choosing another token of the analogical source in the given environment (step 3). After this, the two tokens of the analogical source are compared and their

difference is applied to the analogical target in the known environment (using an algorithm described in Lepage 1998). The result of this operation is the output of the prediction, which, in this case, is [aɪdɪər].

Since steps 1 and 3 are random, the only step that should be discussed in more detail is step 2. The probability of a form S_i being chosen as the analogical source given the analogical target in the known environment, T_j, is shown below:

$$(7) \qquad P(S_i|T_j) = \frac{f_i \eta_{ij}}{\sum\limits_{k \in K} \eta_{jk}},$$

where f_i is the number of tokens the form is exemplified by in the known environment, η is a quantitative measure of similarity, and K is the set of all tokens in the known environment. Since the divisor is constant for all S_i given T_j, the relative probabilities for two different forms are solely determined by their frequency and similarity to T_j. The similarity metric η will not be described in any detail here; the interested reader is referred to Nosofsky (1986). However, it has two important properties, which are worth mentioning here, as they have a significant effect on the outcome of the prediction. First, the tokens were represented by variables in much the same way as in the instance-based models described in Section 4.1. The variables had different weights assigned to them, which means that they influenced the outcome of the prediction to differing degrees. In the simulation presented below, the last sound had the greatest influence. Second, similarity values were attenuated for items that were further away in the feature space, which means that only forms relatively close to the given form had a realistic chance of being chosen as the analogical source (unless they were of very high frequency). Having seen the basic properties of token-based analogy, we can now turn to the simulation itself.

5.1 Simulating the emergence of intrusive-r

The input dataset of the simulation was a set of 1 million tokens randomly chosen from the CE18 corpus, each of them stored in the form presented above (e.g. {[ajdɪə], IDEA, __C}). The transcriptions were modified to reflect a fully rhotic dialect, such as the one spoken in the South of England before the 18th century. To create the conditions for the analogical extensions described above, an additional bias to delete coda-*r* was introduced into the model. The simulation consisted of 30 rounds. In figure 4, I provide a summary of the changes in the dataset, with the dark line indicating the percentage of rhotic productions in prevocalic position within the *r*-less group, the grey line the percentage of rhotic productions in preconsonantal position within the *r*-ful group and the dashed line the percentage of rhotic productions in prevocalic position within the *r*-ful group.

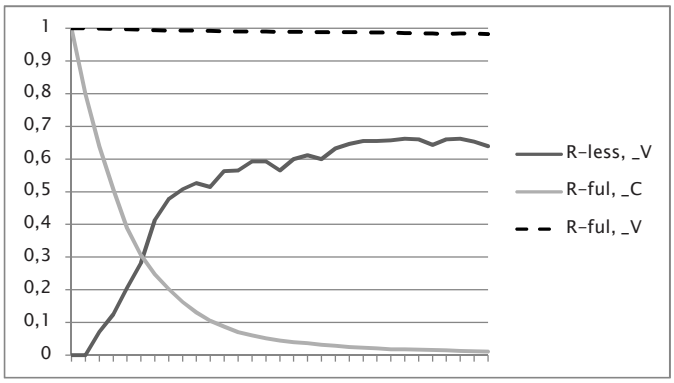

Figure 4: *The loss of rhoticity and the emergence of intrusive-r.*

The changes in Figure 4 clearly indicate that the model presented above is capable of simulating the analogical extensions that led to the emergence of intrusive-*r*. The incidence of *r*-ful productions in prevocalic position rises steadily as the degree of rhoticity decreases, while prevocalic tokens of *r*-ful words do not change their production pattern. The degree of rhoticity and the incidence of intrusive-*r* are negatively correlated; this relationship is strong and significant by a parametric correlation ($r = -0.97$, $p < 0.01$).

The model produces variable results in each round, and the proportion of forms with intrusive-*r* increases only slightly between two generations. However, these small increments slowly add up to create a robust pattern of intrusion. This is perfectly in line with the empirical observations presented in Section 3.2, where I claimed that the development of intrusive-*r* is likely to have been gradual and characterised by a great deal of variation at any given point. Moreover, the extension of the *r*-ful pattern begins well before the model approaches near-categorical non-rhoticity. This, again, corresponds well with the fact that the accents in which intrusive-*r* first appeared are unlikely to have been fully non-rhotic. It should also be pointed out that the variation observed in this model is below the level of lexemes. That is to say, each word develops its own specific – typically non-categorical – distribution of *r*-ful and *r*-less alternants. This is a welcome result in view of the fact that word-specific tendencies have, indeed, been reported for intrusive-*r* (cf. Section 3.2). To sum up, the token-based approach to analogy can provide a highly accurate account of the development of intrusive-*r* in English, which includes a number of factors which previous models of analogy have failed to incorporate.

6 DISCUSSION AND CONCLUSIONS

The preceding sections have demonstrated that the emergence of intrusive-*r* is highly compatible with an analogical approach and that a simulation based on a simple analogical mechanism can produce the pattern of extension that arguably led to the development of this phenomenon in SBE (and possibly other accents as well). It has also been shown that it is possible to account for the gradualness of the development of intrusive-*r* in the same analogical framework if the units over which analogy operates are tokens rather than types. The empirical coverage of this approach and the correctness of its predictions suggest that it is justified to assume that the historical source of intrusive-*r* is analogy.

The implications of these results for synchronic analyses of intrusive-*r* are clear: as it is possible to account for both the history and the present behaviour of intrusive-*r* within a diachronic framework, there is no need for an explanation in purely synchronic terms. Of course, this does not mean that we can dispense with synchronic models altogether: we still have to account for speakers' detailed knowledge of their phonology. As a matter of fact, the token-based analogical model described in Sections 4.2 and 5 makes a number of important assumptions about the nature of this knowledge. Token-based analogy is based on a view of the grammar in which speakers have access to individual instances or exemplars of words and have the ability to make analogical inferences on the basis of these instances. I do not intend to claim that all of phonology boils down to exemplars and analogy; however, it is clear that the present account requires at least these two concepts to be part of the synchronic apparatus of a speaker and can go far in accounting for the phonological facts related to intrusive-*r* without using any additional theoretical machinery.

On a more general note, this account of intrusive-*r* shows that an exemplar-based approach is not necessarily restricted to accounts of phonetic variation, although this is the area where such models have been applied most successfully (cf. Johnson 1997; Pierrehumbert 2001, 2002, 2003). By introducing a simple analogical mechanism and allowing for positive feedback, even relatively weak and variable tendencies can lead to robust patterns. The extensions produced within a single generation in the simulation are sporadic and unsystematic, but the final pattern is highly systematic. Thus, while variability is an important component of the model, it can also deal with systematic alternations that are traditionally considered part of phonology and have not previously been successfully accounted for in exemplar-based approaches.

REFERENCES

Albright, Adam. 2009. Modeling analogy as probabilistic grammar. In James P. Blevins & Juliette Blevins (eds.), *Analogy in grammar: Form and acquisition*, 185–213. Oxford: Oxford University Press.

Albright, Adam & Bruce Hayes. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90(2). 119–161.

Baayen, R. Harald, Richard Piepenbrock & Léon Gulikers. 1995. The CELEX Lexical Database (release 2). University of Pennsylvania, Philadelphia, PA: Linguistic Data Consortium.

Baković, Eric. 1999. Deletion, insertion and symmetrical identity. Ms. Harvard University.

Bauer, Laurie. 1984. Linking /r/ in RP: some facts. *Journal of the International Phonetic Association* 14. 74–79.

Bermúdez-Otero, Ricardo. 2005. The history of English intrusive liquids: using the present to ascertain the past. Handout of paper presented to the Department of Linguistics and English Language, University of Manchester, 24 May 2005. URL: www.bermudez-otero.com/intrusion.pdf.

Blevins, Juliette. 1997. Rules in Optimality Theory: Two Case Studies. In Iggy Roca (ed.), *Derivations and constraints in phonology*, 227–260. Oxford: Oxford University Press.

Brighton, Henry. 2003. Simplicity as a driving force in linguistic evolution. Ph.D. dissertation, University of Edinburgh.

Bybee, Joan L. 2001. *Phonology and language use*. Cambridge: Cambridge University Press.

Chandler, Steve. 2002. Skousen's analogical approach as an exemplar-based model of categorization. In Royal Skousen, Deryl Lonsdale & Dilworth B. Parkinson (eds.), *Analogical modeling*, 51–105. Amsterdam: John Benjamins.

Daelemans, Walter, Jakub Zavrel, Ko van der Sloot & Antal van den Bosch. 2007. TiMBL: Tilburg Memory Based Learner, version 6.1, Reference Guide. ILK Research Group Technical Report Series no. 07-07.

Foulkes, Paul. 1998. English [r]-sandhi: a sociolinguistic perspective. *Leeds Working Papers in Linguistics & Phonetics* 6. 18–39.

Gick, Bryan. 1999. A gesture-based account of intrusive consonants in English. *Phonology* 16. 24–54.

Gick, Bryan. 2002. The American intrusive L. *American Speech* 77. 167–183.

Gimson, Alfred Charles. 1980. *An introduction to the pronunciation of English*. London: Arnold.

Halle, Morris & William J. Idsardi. 1997. r, hypercorrection and the elsewhere condition. In Iggy Roca (ed.), *Derivations and constraints in phonology*, 331–348. Oxford: Oxford University Pres.

Harris, John. 1994. *English sound structure*. Oxford & Cambridge, Mass.: Blackwell.

Hay, Jennifer & Margaret MacLagan. to appear. Social and phonetic conditioners on the frequency and degree of 'intrusive /r/' in New Zealand English. In Dennis Preston & Nancy Niedzielski (eds.), *A reader in sociophonetics*, Berlin: Mouton de Gruyter.

Hay, Jennifer & Andrea Sudbury. 2005. How rhoticity became /r/-sandhi? *Language* 81. 799–823.

Hock, Hans Henrich. 2003. Analogical change. In Brian D. Joseph & Richard D. Janda (eds.), *The handbook of historical linguistics*, 441–460. Oxford: Blackwell.

Jespersen, Otto. 1909. *A modern English grammar on historical principles*. London: George Allen & Unwin.

Johnson, Keith. 1997. Speech perception without speaker normalization: An exemplar model. In Keith Johnson & John W. Mullennix (eds.), *Talker variability in speech processing*, 145–165. San Diego: Academic Press.

Jones, Charles. 1989. *A history of English phonology*. London: Longman.

Jones, Daniel. 1964. *An outline of English phonetics*. Cambridge: Heffer and Sons Ltd 4th edn.

Kirby, Simon, Mike Dowman & Thomas L. Griffiths. 2007. Innateness and culture in the evolution of language. *Proceedings of the National Academy of Sciences* 104(12). 5241–5245.

Lass, Roger. 2000. Phonology and morphology. In Norman Blake (ed.), *The Cambridge history of the English language, Vol. III 1476-1776*, 23–155. Cambridge: Cambridge University Press.

Lepage, Yves. 1998. Solving analogies on words: an algorithm. In *Proceedings of COLING-ACL'98, Montreal, August 1998, vol. 1*, 717–722.

MacMahon, M. K. C. 1998. Phonology. In Suzanne Romaine (ed.), *The Cambridge history of the English language, Vol. IV 1776-1997*, 375–535. Cambridge: CUP.

McCarthy, John J. 1991. Synchronic rule inversion. In L. Sutton, C. Johnson & R. Shields (eds.), *Proceedings of the Seventeenth Annual Meeting of the Berkeley Linguistics Society*, 192–207. Berkeley, CA.: Berkeley Linguistics Society.

McCarthy, John J. 1993. A case of surface constraint violation. *Canadian Journal of Linguistics* 38. 169–195.

McDavid, Raven I. 1958. The dialects of American English. In W. Nelson Francis (ed.), *The structure of American English*, New York: The Ronald Press Company.

McMahon, April. 2000. *Lexical phonology and the history of English*. Cambridge, UK: Cambridge University Press.

Nosofsky, Robert M. 1986. Attention, similarity and the identification-categorization relationship. *Journal of Experimental Psychology: General* 115. 39–57.

Nosofsky, Robert M. 1988. Similarity, frequency and category representations. *Journal of Experimental Psychology: Learning Memory and Cognition* 14. 54–65.

Oudeyer, Pierre-Yves. 2006. *Self-organization in the evolution of speech*. Oxford: Oxford University Press.

Pierrehumbert, Janet B. 2001. Exemplar dynamics: Word frequency, lenition, and contrast. In Joan L. Bybee & Paul Hopper (eds.), *Frequency effects and the emergence of lexical structure*, 137–157. Amsterdam: John Benjamins.

Pierrehumbert, Janet B. 2002. Word-specific phonetics. In Carlos Gussenhoven & N. Warner (eds.), *Laboratory phonology, Vol. VII*, Berlin: Mouton de Gruyter.

Pierrehumbert, Janet B. 2003. Phonetic diversity, statistical learning and acquisition of phonology. *Language and Speech* 46. 115–154.

Rumelhart, David E. & James L. McClelland. 1986. *Parallel distributed processing*. Cambride, MA: MIT Press.

Sheridan, Thomas. 1762/1803. *A course of lectures on elocution. 2nd American edition.* London: Strahan.

Skousen, Royal. 1989. *Analogical modeling of language*. Dordrecht: Kluwer Academic Publisher.

Skousen, Royal, Deryl Lonsdale & Dilworth B. Parkinson. 2002. *Analogical modeling*. Amsterdam: John Benjamins.

Sóskuthy, Márton. 2009. Why *r*? an alternative look at intrusive-*r* in English. M.A. thesis, Eötvös Loránd University, Budapest.

Uffmann, Christian. 2007. Intrusive [r] and optimal epenthetic consonants. *Language Sciences* 29. 451–476.

Wedel, Andrew B. 2004. Self-organization and categorical behavior in phonology. Ph.D. dissertation, University of California at Santa Cruz.

Wedel, Andrew B. 2007. Feedback and regularity in the lexicon. *Phonology* 24. 147–185.

Wells, John. 1982. *Accents of English. 3 Volumes*. Cambridge: Cambridge University Press.

NOTES

[1]Some authors prefer the term postvocalic *r*; I will, however, avoid this term as it is misleading and inaccurate: the word-medial *r* in words such as *ferry* and *zero* is also postvocalic, but it is never involved in *r*∼zero alternations.

[2]The change that I refer to as Pre-*r* Breaking is often analysed as consisting of two separate changes, Pre-*r* Breaking and Pre-*r* Laxing (Wells 1982; McMahon 2000); since Pre-*r* Laxing did not affect the schwa off-glide of centring diphthongs, this distinction is immaterial to the present discussion.

[3]The GCM, which was developed in the field of experimental psychology, is not traditionally grouped together with models like TiMBL and AM, which have their roots in machine learning and linguistics. However, I see no principled reason for making a sharp distinction between the GCM and the other two models apart from their different historical background, as they all rely on stored instances for making predictions and they have a similar underlying architecture (see also Chandler 2002).

[4]This is because only types ending in [ə, ɔː, ɑː, ɜː] ever follow an *r*∼zero pattern, which makes the last sound of the word a good predictor of *r*-fulness. This is also corroborated by TiMBL, which can automatically calculate the optimal weight for each variable, and which chooses the last sound as the most informative feature when given a dataset like the one in Table 3.

[5]Admittedly, this is a relatively crude approach. It is not entirely clear why the model should use analogy to predict a new pronunciation for each form when it could simply retrieve the stored form. Many analogical models (e.g. Bybee 2001; Wedel 2007) assume that analogy only operates in a minority of cases (e.g. when the pronunciation of a given form has become unavailable due to memory decay or other factors) and productions are otherwise faithful to the dataset. However, this approach cannot be straightforwardly extended to a token-based model. If the input dataset is sufficiently large, even infrequent words will be relatively strongly represented. In such a situation, speakers will usually be able to access at least a few tokens for all but the least frequent words (even at relatively high rates of forgetting), and extensions will hardly ever take place.