6 Salience in sociophonetics — a case study of Hungarian hiatus resolution

Péter Rácz
Universität Freiburg

In sociolinguistics, salience is commonly interpreted as a trait that renders a variable more apparent to language users. This paper offers an empirical definition of salience, based on the probability distributions of the realisations of a variable across different dialects. The example used to illustrate this operationalisation of salience is hiatus resolution in Hungarian. I will show that a likely source of salience is that the difference between the probability distributions of its realisations across dialects causes it to be more unfamiliar, and, consequently, surprising for the speakers of one dialect *vis à vis* another one. The method advocated here is most applicable to phonological variables, though adaptation to other language domains is not impossible.

Section 1 overviews the concept of salience in sociolinguistics and introduces its interpretation which is espoused in the rest of the paper. Section 2 discusses Hungarian hiatus resolution, Section 2.2 looks at its social evaluation, based on an attitude test performed with native speakers, Section 2.3 investigates the potential relationship between speaker sensitivity to hiatus resolution and its distributions in language use. Section 3 provides brief conclusions.

1 Salience in sociolinguistics

Salience in sociolinguistics refers to a property or set of properties that cause a language variable to be more prominent, more conspicuous to the language users. The concept is most readily interpreted in the context of the dichotomy between *indicators* and *markers*, introduced by Labov (1972). These concepts are used predominantly when talking about phonological variables. Indeed, salience as such is almost always discussed in the context of phonology — it will not be otherwise in the present paper.

Indicators are variables which vary with social stratification, but have no social interpretation. If we have a standard and a substandard dialect, an indicator variable will be realised differently in the two. Yet, substandard speakers will not try to use the standard realisation when approximating the standard dialect, and this will not be noticed by the standard speakers. That is, indicators do not show style-shifting, and their use by speakers does not invoke value judgements from the members of the language community. They are not subjects of naïve linguistic awareness either. One example is [aː] in Norwich (Trudgill 1986). This vowel is more fronted than the standard variety, but the speakers seem to be unaware of this difference.

Markers are variables which correlate with sociolinguistic identity. If a marked realisation attaches to a substandard dialect, speakers will try to avoid it in more formal style settings and will regard its use as *base* or *erroneous*. An example for a marker could be the Northern [ɑ] (Wells 1982). In the North of England, this sound is restricted to a set environments indicated by a following <r> in the orthography (e.g. *carton*, *bar*). In words like *dance*, *fast*, a fronted [a] is used instead. This realisation of the variable is a strong marker of Northern speech, and Northerners will try to avoid it if conforming to the Southern standard.

Labov et al. (2009) show that if listeners identify a marker realisation as low prestige, this will affect their judgement of speech input even when the realisation is relatively rare in the input. This suggests that the 'detection' and evaluation of markers is independent of the frequency of realisations. (Stuart-Smith 2003 makes a similar observation on Urban Scots in Glasgow.)

The concept of *salience* is discussed, among others, by Trudgill (1986) and Kerswill and Williams (2002). In the interest of brevity, I will not explore these papers in detail. Essentially, they argue that one possible interpretation of salience is to regard it as a cognitive-perceptual property that separates markers from indicators. If salience is, in fact, a property that language users rely on to tell apart indicators and markers, it can have two possible sources. We can either attribute it to speaker dynamics, that is, the organisation of the social space in which language is used, or to a special characteristic that salient markers share but non-salient indicators do not.

The first possibility means that salience is mandated by the language community. That is, any linguistic variable could theoretically be chosen to mark social indexation, independent of the variable's properties. This is the view embraced in Labov (1972). All variables start out as indicators, and later become markers, when the linguistic change gains enough momentum to be noticed by the community, and, as a result, become a vessel of social indexation. This view can be inferred from Labov's Martha's Vineyard study. According to the study, the local residents at Martha's Vineyard picked up on a shift of the realisations of the diphthongs [aw] and [ay] to separate themselves from the summer residents. The small difference between the local and the New England dialect became amplified to mark social identity. At the beginning of this phase, the diphthongs are only indicators of this difference, later, as they start to be used in asserting the local identity, they become markers.

Labov (1994) discards this simple approach to the relationship of indicators and markers, pointing to the fact that some variables never seem to become markers at all. If the basis of salience is not only social dynamics, one ought to find a general, perceptual frame, that prefers some variables to others. Both Trudgill (1986) and Kerswill and Williams (2002) point to general cognitive capacities as possible sources of sociolinguistic salience. The supposition is reasonable: some variables might be picked up because they are more highlighted in the course of acquisition or auditory perception.

The nature of the perceptual and cognitive properties that go with salience are not clearly established by any of these authors. In this paper, I will claim that the salience of a variable comes from its patterning in use: *some variables are more surprising for speakers of a different dialect, and, consequently, carry social indexation easier*.

Assuming a strictly segmental approach, surprisal can be measured explicitly relying on the notion of transitional probabilities (TP-s). The transitional probability of a segment Y following segment X is the chance that we find Y immediately following X in a given corpus (cf. Table 1).

$$p(\mathrm{Y}|\mathrm{X}) = \frac{\text{likelihood of pair XY}}{\text{likelihood of X}}$$

Figure 1: *Probability of Y following X*

The use of TP-s in linguistics was first suggested by Harris (1955), who proposed that a field linguist can rely on them when transcribing an unknown language. Since the ordering of segments within a word is constrained, but (almost) any pair of segments can occur with an intervening word boundary, some patterns (the ones permitted in words) will occur more frequently than others (the ones only occurring at boundaries). A low TP, in turn, hints at a word boundary.

A large body of research suggests that not only field linguists but also language users are capable of using such statistical information in locating word boundaries (Jusczyk et al. 1994; Saffran et al. 1996; Cairns et al. 1997; Pierrehumbert 2003; Hay 2000). The question of how listeners find word boundaries in the speech signal is not uncontroversial. Still, probability-based statistical learning seems to play a prominent role in it, both in the case of infant and adult learners. Though listeners certainly rely on other distributional cues, such as word stress, pauses in the signal, or simply the recognition of words previously heard in isolation, transitional probability between the segments remains the most abundant and reliable cue (Saffran et al. 1996; Jusczyk et al. 1999).

If we accept the role of transitional probabilities in segmenting the speech signal into words, it is straightforward to assume that this type of statistical information is, to an extent, available for language users. Consequently, a variable realisation that strongly alters the TP-s will be salient for language users. In this reading, salience comes into play when comparing two dialects or idiolects, in which the distributions of a particular variable realisation are notably different. In the following section, I give an example on the relationship of salience and low transitional probabilities, hiatus resolution in Hungarian.

2 Hiatus resolution in Hungarian

In this section, I discuss two types of hiatus resolution in Hungarian. I give the results of an attitude test which support that one of these types is salient for the language users, while the other one is not. Finally, I give an approximation of the transitional probabilities of the two types, linking the difference in salience to the difference in TP-s.

2.1 Types of hiatus resolution

Educated Colloquial Hungarian (ECH) has hiatus resolution in two distinct environments. The first sort is obligatory and non-salient in the standard, while the second occurs to a much smaller extent — if at all — and is subject to variation, as well as distinctly salient (Siptár and Törkenczy 2000). It is present in many other Hungarian dialects, but that is beyond our scrutiny.

The phonetics and phonology of hiatus resolution has been extensively covered (Kálmán and Rebrus 2010; Siptár and Törkenczy 2000; Siptár 2003), but its social evaluation has been scarcely discussed in any depth. Siptár and Törkenczy (2000), whose description I mainly rely on, only mention the issue in *passim*. The basic state of affairs is as follows: Hungarian has lexical and post-lexical hiatus resolution. Our focus is post-lexical hiatus resolution (cf. Table 1).[1] It occurs obligatorily in vowel clusters containing [i] and it is quite common in clusters containing [eː]. The inserted segment is the glide [j]. These are the two close front vowels of Hungarian. The close [i], like all Hungarian vowels, has a long pair [iː]. However, in ECH, the realisation of close vowels is subject to variation, and they generally show

[1] The morphological make-up of the hiatuses is not relevant for the present discussion: vowel clusters with [i] have obligatory hiatus filling without respect to the presence of a boundary, while clusters with [eː] and [ɛ] are relatively rare in mono-morphemic words, dwarfing the effect of this factor in a corpus study.

a tendency to be shortened in all positions except the initial syllable. This bears no importance on the present discussion. The short pair of [eː] is also different qualitatively, and it is realised as an open mid [ɛ].

Vowel clusters containing [ɛ] but *not* [i] or [eː] also show hiatus resolution, it is, however, realised less frequently, and, in any case, is subject to variation. There is no data on the extent of hiatus resolution in the three environments, namely, in clusters involving [i], [eː], and [ɛ], respectively, but it is commonly assumed that, in Conservative ECH, hiatus resolution is obligatory in the first, variable in the second, and avoided in the third environment. However, as Siptár and Törkenczy point out, Innovative ECH has hiatus resolution in the third environment, although to a limited extent.

It has to be noted here that hiatus resolution with [eː] and [ɛ] is not only variable in its occurrence but also in its extent. Certain realisations are likely to be, for instance, longer in duration than others. This is an issue that I will not take into consideration here, but which is certainly relevant for further study.

fiú	[fijuː]	'boy'
női	[nøːji]	'female'
ráér	%[raːjeːr]	'to be at leisure'
büféasztal	%[byfeːjɒstɒl]	'buffet table'
tea	%%[tɛjɒ]	'tea'
beakad	%%[bɛjɒkɒd]	'gets stuck'

Table 1: *Post-lexical hiatus resolution in Hungarian*

I will argue that, compared to the first two environments, the third one is salient for the speakers of ECH. I am going to underpin my argument by discussing the results of an attitude test on the perception of this variable. The test provides empirical evidence on the pattern's salience.

2.2 Salience and hiatus resolution

Methods

The test included ten ECH speakers, five female, five male, with a mean age of 22. Eight were from Budapest, and two from the surrounding Pest county. The participants listened to a recording of 30 sentence pairs, 10 with vowel plus [ɛ] clusters, 10 with [ɛ] plus vowel clusters, and ten control sentences, featuring V[i]/[i]V and

V[eː]/[eː]V equally. The recordings were mono audio waveform files, sampled at 44100Hz. The participants were not paid for the experiment.

The test sentences were read by a trained phonetician, also a native speaker of ECH, once with a hiatus filler [j], once without one. The control sentences were all read with a realised hiatus filler. The background information given to participants was that a Hungarian male in his twenties is looking for their help in general linguistic and stylistic issues, as he is going to a job interview in Budapest and is unsure about the quality of his Hungarian. The participants had to evaluate the sentence pairs (with the implication that the sentences are different) on a Likert-scale from 1 to 10, depending on whether they found the first or the second sentence better (or they were unsure, etc.). The participants listened to the pairs in a random order, both in the sense that the order of pairs was randomised and that the order of the *marked* sentence (the one with hiatus resolution) and the *unmarked* sentence (the one without it) was randomised: half the pairs had the marked sentence first. The listening test was followed by a small discussion with the participants.

The experiment has two conditions: (i) whether the marked sentence comes first or second and (ii) whether the judgements on the test sentences differ from judgements on the control sentences. The hypothesis is that hiatus resolution with [ɛ] is a salient variable that will be rejected, whereas hiatus resolution with [i] and [eː] elicits no listener attitudes. This should show up in condition (i) as a larger score on the scale if the first sentence is marked in the pair and *vice versa*, and in (ii) as a score more divergent from the mean in the case of test sentences versus control sentences, as participants are not expected to show explicit preference for any sentence in a pair of control sentences.

Results

The results show a strong preference for the unmarked pattern in condition (i) and more divergent scores in condition (ii), which confirms the hypothesis that hiatus resolution in [ɛ]+V clusters is rejected, hence, salient for ECH speakers.

The results were weighted between participants. For condition (i), the resulting scores were modified in such a way that a higher score means a preference for the *marked* pattern. Condition (ii) is needed in the first place because condition (i) relies on the order of sentences within the pairs. Therefore, the results can be influenced not only by which sentence was marked, but also by the order itself: if the first sentence is marked, chances are, people become more aware of it. Since condition (ii) compares all the test sentences with the control sentences, the problem of ordering disappears.

The scores are shown in Figure 2 for the first condition and Figure 3 for the second one. In Figure 2, the first column is the control, the second is where the marked sentence came first in the pair, and the third is when it came second. Higher scores indicate a stronger overall preference for the second sentence. As can be seen, if the first sentence is marked, scores are higher, and if the second sentence is marked, they are lower than in the control case.

Figure 3 compares control and test sentences. Higher scores generally indicate a stronger overall preference for the marked sentence. Of course, the control pairs did not have a marked sentence. Scores are again weighted, which resulted in a higher score than the expected 0 for the control case. What is visible, however, is the existence of a deviation from the mean in the test case, which is absent in the control case. That is, while there was a preference for one member in the pair in the test case, this preference was absent in the control case.

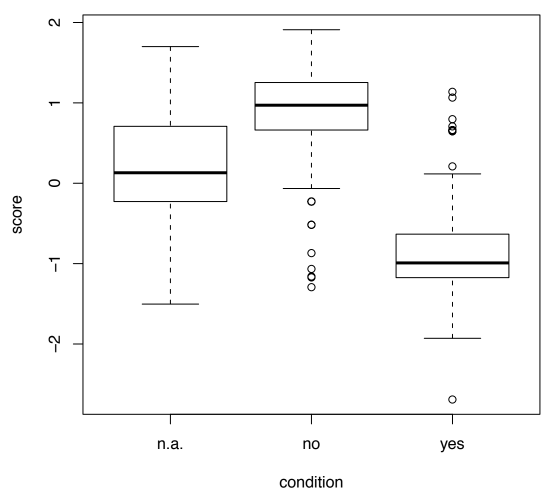

Figure 2: *Weighted scores for condition (i)*

The significant difference between answer rates, visible on the plots, is supported by a fitted linear mixed model for both conditions, with speakers (`subject`) and read sentence pairs (`sentence.id`) as a random effect. The lme4 package (Bates 2005), implemented in R (R Development Core Team 2009), was used for the mixed-effects modelling. The summary of the model for condition (i) is in (1) and the summary of the model for condition (ii) is in (2). We can see that in

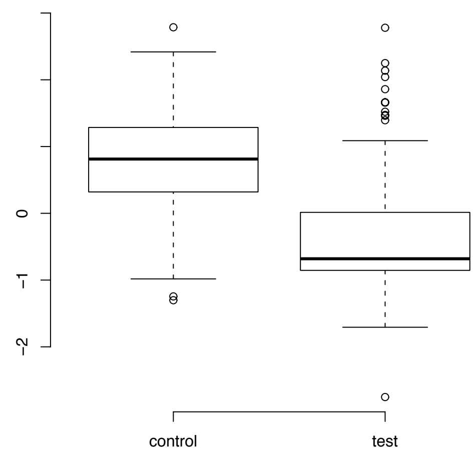

Figure 3: *Weighted scores for condition (ii)*

condition (i), whether the first or the second sentence was marked in the pair (`2nd_markedno/yes`) is a significant predictor of the weighted score (`score`), with a strong correlation of the fixed effects. In condition (ii), whether the sentence was a condition or a test sentence (`conditiontest`) is again a strong predictor of the weighted score (`alt.score`), with an observable strong correlation again.

(1) Summary of the model for condition (i)

```
Linear mixed model fit by REML
Formula: score ~ 2nd_marked + (1 | subject) + (1 | sentence.id)
   Data: dat
   AIC   BIC logLik deviance REMLdev
 749.9 772.1 -368.9    730.1   737.9
Random effects:
 Groups      Name        Variance Std.Dev.
 sentence.id (Intercept) 0.054274 0.23297
 subject          (Intercept) 0.000000 0.00000
 Residual                0.633746 0.79608
Number of obs: 300, groups: sentence.id, 30; subject, 10
Fixed effects:
                  Estimate Std. Error t value
(Intercept)         0.7434     0.1085   6.854
```

```
2nd_markedno   -0.9542     0.1576  -6.055
2nd_markedyes  -1.2466     0.1499  -8.318

Correlation of Fixed Effects:
             (Intr) X2nd_s_mrkdn
2nd_mrkdn -0.688
2nd_mrkdy -0.724  0.498
```

(2) Summary of the model for condition (ii)

```
Linear mixed model fit by REML
Formula: alt.score ~ condition + (1 | subject) + (1 | sentence.id)
   Data: dat
   AIC BIC logLik deviance REMLdev
 749.5 768 -369.7    733.7   739.5
Random effects:
 Groups      Name        Variance Std.Dev.
 sentence.id (Intercept) 0.065188 0.25532
 subject          (Intercept) 0.000000 0.00000
 Residual                0.633746 0.79608
Number of obs: 300, groups: sentence.id, 30; subject, 10
Fixed effects:
              Estimate Std. Error t value
(Intercept)     0.7434     0.1134   6.556
conditiontest  -1.1151     0.1389  -8.030

Correlation of Fixed Effects:
            (Intr)
conditintst -0.816
```

In sum, the results confirm speaker awareness of the Innovative hiatus resolution pattern *vis-à-vis* the Conservative pattern. For a Hungarian linguist, this is hardly surprising, as the pattern is overtly discussed, and some forms like *teja* are used playfully by speakers who otherwise eschew Innovative hiatus resolution. It is, however, important to stress that the pattern's social evaluation was not empirically tested before.

2.3 Salience and TP-s in Hungarian hiatus resolution

This section looks at the correlation between salience and transitional probabilities in Hungarian hiatus resolution. The procedure is the following: I take a written corpus of Hungarian, and modify it in such a way that it includes [j]-s resulting from standard hiatus resolution, not marked in the orthography. I extract the frequencies of [ij] and [ji] clusters in order to gain the transitional probabilities (TP-s) of [j]||[i] and [i]||[j] in the corpus. I also extract the frequencies of [ɛj] and [jɛ] to gain the TP-s of [j]||[ɛ] and [ɛ]||[j].

The hypothesis is that pairs of [j] and [i] are much more frequent than pairs of [j] and [ɛ], that is, the TP-s of [j]||[i] and [i]||[j] are larger than those of [j]||[ɛ] and [ɛ]||[j], respectively. Consequently, the former are more familiar to the listeners, so when these occur as a result of hiatus resolution, the pattern is not salient. In comparison, the latter are much less familiar, so when these result from hiatus resolution, the pattern becomes salient. Again, the key point is that clusters of, for instance, [ɛj] are not illicit and occur in Hungarian, their salience in hiatus resolution comes from the frequency difference.

The data are drawn from the Hungarian Webcorpus (Halácsy et al. 2004), a corpus of 1.48 billion words from 18 million pages downloaded from the .hu Internet domain, which gives the best representation of written language, and is the most faithful corpus of present-day Hungarian. A sample of 17 million words was used to establish TP-s. Hungarian orthography is relatively consistent, at least when it comes to the representation of [j], [i], and [ɛ]. It does not mark hiatus resolution, so I inserted [j]-s into vowel clusters including [i]. This step is valid inasmuch as hiatus resolution is obligatory in these clusters. It assumes, however, that hiatus filler [j]-s are equal to contrastive [j]-s in the language. This assumption is supported by authors like Kálmán and Rebrus (2010), who argue that the intrusive segment in hiatus resolution is phonologically equal to the one in the possessive. This, in turn, means that all intervocalic [j]-s are interpreted equally, as the possessive [j] is virtually indistinguishable from the contrastive 'lexical' one. An example to this is given in Table 2.

Environment	Process
zoknija [zoknijɒ] 'sock-POSS3SG'	Possessive suffixation
szoknia [soknijɒ] 'accustom-INF3SG'	Hiatus resolution before infinitive *-a*
kijavít [kijɒviːt] 'fix-3SG'	Contrastive [ijɒ] sequence

Table 2: *[j] in Hungarian*

Even if we take the ontology of [j] as granted, the analysis has to cope with another difficulty, the lack of reliable data on hiatus resolution. In Conservative ECH, it is agreed to be obligatory in vowel clusters with [i] and variable in clusters containing [eː]. There are no estimates on Innovative ECH. In order to tackle the scarcity of the data, I take up the approach of looking at transitional probabilities in one dialect instead of comparing two.

This dialect, Conservative Educated Colloquial Hungarian, is assumed to be represented by the Webcorpus. It has obligatory hiatus resolution with [i], but has none with [ɛ] (since it is conservative). Nonetheless, it also has 'lexical' instances of vowel clusters with [i]/[ɛ] and [j]. This is illustrated in Table 3.

Lexical sequence	Hiatus
kijárat [kijaːrɒt] 'exit'	*kiárad* [kijaːrɒd] 'flow-3SG'
Tejút [tɛjuːt] 'Milky Way'	*szemleút* [sɛmlɛuːt] 'field trip'

Table 3: *Conservative ECH*

With these presumptions, we can look at the frequency differences of non-salient and salient hiatus resolution in the corpus. The frequency of the relevant string in the corpus is given in Table 4. (Both word-internal clusters and clusters including a word boundary were included.) The TP-s are given in Table 5 (numbers are rounded to the third decimal place).

String	Frequency
ɛj	103024
jɛ	230857
ij	480943
ji	391069
i	4424703
j	2367677
ɛ	10892098

Table 4: *String frequencies in the corpus sample*

The results show that there is a frequency difference of one order of magnitude between the TP of [j] following [i] versus that of [j] following [ɛ]. There is no difference, however, when we look at the pattern the other way around, that is, between the TP of [i] following [j] versus [ɛ] following [j]. This asymmetry can be probably blamed on the possessive suffix, which is *-jɛ* after vowel-final front vowel stems.

Environment	TP
j|ɛ	0.009
ɛ|j	0.098
j|i	0.109
i|j	0.088

Table 5: *TP-s in the corpus sample*

What the corpus study tells us, then, is that the salience of the innovative hiatus resolution pattern shows a correlation with the relative low frequency of the string [jɛ] (when compared to [ji]), itself part of the realisation of the resolved hiatus. This supports the hypothesis that the salience of the innovative hiatus resolution pattern, confirmed by the attitude study, springs from a difference in transitional probabilities, which difference renders the pattern less familiar to the listeners. (The grapheme <í>, indicating a long vowel that often undergoes shortening in the spoken language, was not included in the counts. All other things being equal, this should not affect the results.)

Two questions should be addressed at this point. First, it has to be stressed that there is no data available on the Innovative ECH dialect assumed here, apart from its existence. The extent of hiatus resolution in vowel clusters with [ɛ], as well as its origin and correlation with innovative realisations of other variables remain subject to a future study. Second, one might argue that if, based on the corpus data, instances of [j] following [ɛ] are unfamiliar, we should expect the salience of any [ɛj] sequence, not just the ones arising through hiatus resolution. To put the question differently: why is *teja* salient, but *bejárat* 'entrance' apparently not? In my view, the difference lies in the patterning of the two types of clusters. Lexical, contrastive clusters show no variation in ECH. In this sense, one cannot talk about conservative and innovative use, or, indeed, about a linguistic variable. The rarity of [ɛj] clusters is relevant where these clusters occur variably.

3 Conclusions

In this paper I looked at salience in sociolinguistics, and defined it as a perceptual property separating indicators from markers, in the sense of Labov (1972). (For a different interpretation of salience, see for instance Trudgill 1986.) The salience of a variable was operationalised as a by-product of the different transitional probabilities of its realisations in two dialects, causing its patterning in dialect B to be surprising or unusual for a speaker of dialect A.

The approach was explicated by a look at hiatus resolution in Standard Colloquial Hungarian. It was argued that speaker awareness of hiatus resolution in [ɛ]+vowel sequences stems from the different distributions of the resulting sequences in Innovative ECH, where this pattern is present, as opposed to Conservative ECH, where it is mostly absent. A vital point to make is that while the relevant sequences can be found in both dialects, their probability of occurrence is different. Speaker awareness of the variable was supported by an attitude test, which gave hitherto lacking empirical support to the salience of the variable. The differences in distributions were calculated based on corpus data. This view of salience favours the study of phonological variables, but it can be extended to account for other variables as well.

Tackling salience in such a way is empirically fruitful as it gives well-defined, testable tools in the investigation of variable behaviour. Though the rigid segmental approach employed implies a perhaps unwanted level of abstractness, and does not permit the investigation of all phonological variables, such as differences in vowel quality, it is a step forward from intuitive formulations on salience and its influence on the social life of linguistic variables.

Bibliography

Bates, D. (2005). Fitting linear models in R using the lme4 package. *R News 5(1)*, 27–30.

Cairns, P., R. Shillcock, N. Chater, and J. Levy (1997). Bootstrapping word boundaries: A bottom-up corpus-based approach to speech segmentation. *Cognitive Psychology 33*, 111–153.

Halácsy, P., A. Kornai, L. Németh, A. Rung, I. Szakadát, and V. Trón (2004). Creating open language resources for Hungarian. In *Proceedings of the 4th international conference on Language Resources and Evaluation (LREC2004)*, pp. 203–210. Lisbon: European Language Resources Association.

Harris, Z. S. (1955). Phoneme to morpheme. *Language 31*, 190–222.

Hay, J. (2000). *Causes and consequences of word structure*. Ph. D. thesis, Northwestern University.

Jusczyk, P. W., D. M. Houston, and M. Newsome (1999). The beginnings of word segmentation in english-learning infants. *Cognitive Psychology 39*, 159–207.

Jusczyk, P. W., P. A. Luce, and J. Charles-Luce (1994). Infants sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language 33*, 630–645.

Kálmán, L. and P. Rebrus (2010). Valóban megmagyarázhatatlanok a magyar infinitívusz toldalékai? [Are the Hungarian infinitive suffixes really inexplicable?]. In *A mai magyar nyelv leírásának újabb módszerei. VII*. Szeged: SZTE.

Kerswill, P. and A. Williams (2002). 'Salience' as an explanatory factor in language change: evidence from dialect levelling in urban England. *Contributions to the sociology of language 86*, 81–110.

Labov, W. (1972). *Sociolinguistic patterns*. University of Pennsylvania Press.

Labov, W. (1994). *Principles of linguistic change: internal factors*. Wiley-Blackwell.

Labov, W., S. Ash, M. Ravindranath, T. Weldon, M. Baranowski, and N. Nagy (2009). Listeners' sensitivity to the frequency of sociolinguistic variables. Ms.

Pierrehumbert, J. (2003). Probabilistic phonology: discrimination and robustness. In *Probabilistic Linguistics*. Cambridge, MA: MIT Press.

R Development Core Team (2009). *R: A Language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.

Saffran, J. R., R. N. Aslin, and E. L. Newport (1996). Statistical learning by 8-month-old infants. *Science 274*, 1926–1928.

Saffran, J. R., E. L. Newport, and R. N. Aslin (1996). Word segmentation: the role of distributional cues. *Journal of Memory and Language 35*, 606–621.

Siptár, P. (2003). Hungarian yod. *Acta Linguistica Hungarica 50*, 457–473.

Siptár, P. and M. Törkenczy (2000). *The phonology of Hungarian*. Oxford: Oxford University Press.

Stuart-Smith, J. (2003). The phonology of modern urban Scots. In *The Edinburgh Companion to Scots*. Edinburgh: Edinburgh University Press.

Trudgill, P. (1986). *Dialects in contact*. Oxford: Blackwell.

Wells, J. C. (1982). *Accents of English*. Cambridge: Cambridge University Press.