Available online at www.sciencedirect.com

ScienceDirect

Lingua 171 (2016) 74–91

Lingua

www.elsevier.com/locate/lingua

Is speech processing influenced by abstract or detailed phonotactic representations? The case of the Obligatory Contour Principle

Natalie Boll-Avetisyan [a,b,*], René Kager [a]

[a] Utrecht University, Trans 10, 3512 BL Utrecht, The Netherlands
[b] University of Potsdam, Department of Linguistics, Karl-Liebknecht-Str. 24-25, 14476 Potsdam, Germany

Received 25 October 2014; received in revised form 7 November 2015; accepted 27 November 2015
Available online 14 January 2016

Abstract

Many languages restrict their lexicons by OCP-Place, a phonotactic constraint against co-occurrences of consonants with shared [place] (e.g., McCarthy, 1986). While many previous studies have suggested that listeners have knowledge of OCP-Place and use this for speech processing, it is less clear whether they make reference to an abstract representation of this constraint. In Dutch, OCP-Place gradiently restricts non-adjacent consonant co-occurrences in the lexicon. Focusing on labial-vowel-labial co-occurrences, we found that there are, however, exceptions from the general effect of OCP-Labial: (A) co-occurrences of identical labials are systematically less restricted than co-occurrences of homorganic labials, and (B) some specific pairs (e.g., /pVp/, /bVv/) occur more often than expected. Setting out to study whether exceptions such as (A) and (B) had an effect on processing, the current study presents an artificial language learning experiment and a reanalysis of Boll-Avetisyan and Kager's (2014) speech segmentation data. Results indicate that Dutch listeners can use both knowledge of phonotactic detail and an abstract constraint OCP-Labial as a cue for speech segmentation. We suggest that whether detailed or abstract representations are drawn on depends on the complexity of processing demands.

Keywords: OCP; Artificial language learning; Speech segmentation; Phonotactics; Identity; Phonotactic probability

1. Introduction

Languages' lexicons are restricted by phonological constraints that determine which adjacent and non-adjacent within-word combinations of sounds are likely, unlikely or even illegal. There is ample evidence that humans have knowledge of phonotactic probabilities in their language input, which influences their wellformedness intuitions about nonwords (e.g., Bailey and Hahn, 2001; Frisch et al., 2000; Scholes, 1966) and lexical access (e.g., Berent et al., 2001b; Vitevitch and Luce, 1998, 1999). Moreover, listeners can use phonotactic cues for speech segmentation (e.g., Boll-Avetisyan and Kager, 2014; McQueen, 1998; Suomi et al., 1997; Vroomen et al., 1998) and benefit from phonotactic knowledge when learning new words (e.g., Boll-Avetisyan, 2012a; Gathercole et al., 1999; Storkel, 2001). In spite of this vast interest in the

* Corresponding author at: University of Potsdam, Department of Linguistics, Karl-Liebknecht-Str. 24-25, 14476 Potsdam, Germany. Tel.: +49 0331 9772374.
E-mail addresses: nboll@uni-potsdam.de (N. Boll-Avetisyan), r.w.j.kager@uu.nl (R. Kager).

http://dx.doi.org/10.1016/j.lingua.2015.11.008
0024-3841/

use of phonotactic knowledge in speech processing and language acquisition in the literature, much less is known of how phonotactic knowledge is represented.

The current study sets out to contribute to the discussion of how exactly phonotactic knowledge is represented in the speakers' minds. In principle, there are three possibilities that have been suggested in the literature: first, phonotactic representations could be detailed. That is, listeners might represent each possible phoneme combination associated with knowledge of its frequency or probability of occurrence — a view classically taken by connectionists (e.g., Plaut et al., 1996). Second, listeners might represent phonotactics by means of abstract constraints in the sense that they abstract over phonological features. For example, if co-occurrences of consonants that are part of a natural class (e.g., all labial consonants) are systematically unattested in a language, listeners would only represent a single abstract constraint, namely *[labial][labial]. This abstract representation would be more economic rather than having many separate specific representations such as */pb/, *pm/, */fm/ et cetera, as it subsumes all specific instances enumerated. This latter view is classically taken by generative phonologists (e.g., Chomsky and Halle, 1968). A third possibility is that knowledge is both abstract and detailed (e.g., Adriaans and Kager, 2010; Albright and Hayes, 2003; Hayes, 1999; Hayes and Wilson, 2008).

The current study investigates whether abstract or detailed knowledge of phonotactics influences speech segmentation. The abstract constraint under test is OCP-Labial, a constraint restricting the co-occurrence of consonants sharing the feature [labial]. In Dutch, effects of OCP-Labial are not absolutely systematic. We focus on two issues: first, we will investigate whether Dutch listeners attribute a special role to co-occurrences of identical labials. Second, we will investigate whether specific pairs of labials that are likely to co-occur in a language, and hence are exceptions from OCP-Labial, will still fall under the effect of the constraint.

1.1. OCP-Place

The *Obligatory Contour Principle* (OCP) is a typologically well-attested constraint on within-word phonological structure requiring co-occurring segments to be featurally non-identical (i.e., to disharmonize). Originally, the OCP was formulated to account for similarity avoidance between tones in West-African languages (Leben, 1980). Later, the OCP was extended to OCP-Place to formally account for co-occurrence constraints against consonants with shared [place] (McCarthy, 1986), as, for example, in Arabic (Greenberg, 1950), where verbs (e.g., *katab-a* 'he wrote', *kutib-a* 'it was written', and *kuttib-a* 'he was made to write') are derived from root morphemes consisting of three consonants (e.g., /k t b/), which, with very rare exceptions, are not homorganic (e.g., */f b m/), or, if root-initially, identical (e.g., */b b k/).

Effects of OCP-Place have first been found in the lexicons of the Semitic language family (e.g., Frisch et al., 2004; Greenberg, 1950; McCarthy, 1985), and were later discovered in many genetically and geographically unrelated languages (e.g., English: Berkley, 1994; Muna: Coetzee and Pater, 2008; Niger-Congo languages: Pozdniakov and Segerer, 2007; Dutch: Kager and Shatzman, 2007). In fact, recent large-scale quantitative typological studies (Graff, 2012; Mayer et al., 2010; Pozdniakov and Segerer, 2007) indicate that virtually every language is restricted by similar place avoidance (c.f. Boll-Avetisyan, 2012b; Graff, 2012: OCP-effects could not be found for the languages Mandarin Chinese, Guarani, Wantoat and Tatar). In light of this very strong typological preference, OCP-Place has been proposed to be a (statistical) universal (e.g., Frisch, 2004; McCarthy, 1986; Pozdniakov and Segerer, 2007). Classically, it is assumed that OCP-Place is a constraint that abstracts over features, i.e., it is represented as *[place][place] (McCarthy, 1986). If sub-classes of consonants with shared place are affected to different degrees, then OCP-Place can be further broken down into constraints such as OCP-Labial (i.e., *[labial][labial]) or OCP-Dorsal (i.e., *[dorsal][dorsal]) (Coetzee and Pater, 2008; McCarthy, 1988).

The majority of languages does not display categorical effects of OCP-Place targeting all homorganic consonant combinations. Rather, most languages display gradient effects of this constraint. This means that pairs of consonants with shared place features can be found in the lexicon; however, they are systematically underattested. That is, pairs such as /pVm/ in English *spam* or /mVb/ in *mob* do occur, albeit significantly less often than expected if non-adjacent consonants co-occurred at random. Frisch et al. (2004) observed that even in the case of Arabic, OCP-effects are graded: forms with roots with a higher overlap in featural identity are less frequent than forms in which roots are less similar. Consequently, they propose a single constraint of similarity avoidance that is quantitatively sensitive to the degree of violation as defined by phonetic feature distance. For Arabic, this is an elegant way to account for the co-occurrence dependencies between consonants that captures both the gradient and abstract aspects of OCP-Place in a single analysis. Other studies, however, have raised doubts that Frisch, Pierrehumbert and Broe's model can be applied to many other languages (e.g., Berent and Shimron, 2003; Coetzee and Pater, 2008; Graff and Jaeger, 2009). Coetzee and Pater, for example, used a model of weighted constraints referring to more specific feature-based sub-classes of consonants (as inspired by, a.o., Padgett, 1995). This approach was more successful in modeling gradient consonant co-occurrence restrictions in Arabic and Muna than Frisch et al.'s similarity metric (Coetzee and Pater, 2008).

In the current study, we address the question of whether the assumption of an abstract representation of OCP-Labial is justified, if the effects of the constraint are gradient. The classical arguments for abstractness are that productivity can best be accounted for if abstract symbolic representations are assumed, and furthermore, they would be beneficial for reducing storage space (Chomsky and Halle, 1968). Furthermore, symbolic representations may facilitate speech processing (e.g., Dupoux et al., 2001a). Connectionists, however, suggest that an inclusion of a grammar component is unnecessary, as the issue of productivity can be explained by associations (e.g., Plaut et al., 1996). Newer proposals integrate the two views in models which combine abstract phonological grammar with a stochastic component (Boersma and Hayes, 2001; Coetzee and Pater, 2008; Frisch et al., 2004; Hayes, 2000; Pierrehumbert, 2003a,b). Such hybrid models leave room for the possibility that listeners generalize over phonotactic distributions if they systematically affect segments that share major class features (e.g., Coleman and Pierrehumbert, 1998; Frisch and Zawaydeh, 2001; Hay et al., 2004; Pierrehumbert, 2003a,b).

1.2. Psychological reality of the OCP

Over the past years, an increasing number of studies has provided experimental evidence for the psychological reality of the OCP. However, there are still some unsolved issues with respect to the exact nature of how identity and/or similarity avoidance is mentally represented. In particular, previous studies provide an inconsistent picture regarding the question of how abstract effects of OCP may be. Some of this literature has concentrated on the role of the OCP restricting the co-occurrence of identical consonants in root morphemes in Modern Hebrew. These studies showed that listeners judge nonwords that contain roots with repetitions of a consonant at the beginning (e.g., /ssm/), which is illegal in Hebrew, as less well-formed than nonwords that contain possible roots (e.g., /mss/ or /psm/; Berent and Shimron, 1997; Coetzee, 2009). Moreover, the OCP seems to affect online speech processing: in lexical decision tasks, nonwords are rejected faster by native listeners of Hebrew if they violate OCP than if they do not (Berent et al., 2001b).

Berent and Shimron (1997) drew the conclusion that Hebrew listeners had abstract knowledge of a markedness constraint against identity between the first and second consonant in a root. However, in their study, items were not controlled for lexical statistics, which, as noted by Frisch and Zawaydeh (2001), could also explain the results. Controlling for this confound, later studies still found effects of identity (e.g., Berent et al., 2001a), suggesting the effects may result from abstract knowledge. Moreover, Berent et al. (2002) showed that native speakers of Hebrew transfer their well-formedness intuitions to novel roots made up of phonemes that are not part of the Hebrew phoneme inventory, which indicates that they generalize over the restriction against root-initial geminates. This suggests an abstract representation.

Other studies have focussed on the effects of OCP-PLACEon homorganic consonant co-occurrences. These studies also found effects of the constraint on wellformedness intuitions (Arabic: Frisch and Zawaydeh, 2001, Dutch: Boll-Avetisyan and Kager, 2014) and lexical decisions (Arabic: Frisch and Zawaydeh, 2001; Dutch: Kager and Shatzman, 2007). Furthermore, OCP-PLACE may play a role during speech segmentation: when learning an artificial language, Dutch adults preferred assuming word boundaries between two consecutive labials rather than between labial-coronal or coronal-labial pairs (Boll-Avetisyan and Kager, 2014). The stimuli in Frisch and Zawaydeh's (2001) as well as Kager and Shatzman's (2007) and Boll-Avetisyan and Kager's (2014) study were highly controlled for lexical statistics, which may indicate an abstract mental representation of OCP-PLACE.

A relevant question to ask is how generalizable such results are if in a language's lexicon OCP-effects are gradient, probabilistic, or even irregular. Potentially, there are two general types of irregularities to consider: (a) systematic exceptions, and (b) unsystematic exceptions.

(a) *Systematic exceptions*: A language might generally restrict the co-occurrence of consonants with shared place, but systematically allow for exceptions. An example of systematic exceptions is constraints on the co-occurrence of identical segments, which often seem to be systematically exempt from general effects of OCP-PLACE. OCP-PLACE then rather holds for homorganic but non-identical consonants. This phenomenon, referred to as an ''escape hatch from OCP-PLACE'' (Yip, 1989), can be encountered in many languages (e.g., Muna: Coetzee and Pater, 2008; English: Frisch, 1996; Aymara, Hausa, and other languages: MacEachern, 1999; Javanese: Yip, 1989). Also, it seems that there is an identity effect on production: native speakers of Cochabamba Quechua made fewer errors when repeating nonwords in which non-adjacent ejective consonants were identical than when they were near-identical, although both patterns are equally unattested in Cochabamba Quechua (Gallagher, 2014). For this reason, it has been proposed that restrictions holding on identical consonants are distinct from those holding on nonidentical homorganic consonants (MacEachern, 1999). In a recent paper, Berent et al. (2012) propose that identical segment co-occurrence restrictions are represented as a grammatical function requiring identity between variables: ''*XX, where X = a consonant'' (Berent et al., 2012, p. 101). They observe that restrictions such as *[labial][labial] can formally not be distinguished from restrictions such as *[coronal][labial], which do not involve identity or similarity between elements. This latter type of constraint is also abstract in the sense that it generalizes over features. However, it could not be expressed as an operation on variables. This view that constraints on identity are represented as constraints on variables receives support from an artificial language learning study, in which participants were able to learn that constraints on identical consonants are distinct from constraints on homorganic consonants (Gallagher, 2013).

Systematic exceptions from OCP have been addressed by a few studies. For example, Berent and Shimron (2003) found that native speakers of Hebrew judged nonwords with identical root-final consonants to be more wellformed than nonwords containing homorganic root-final consonants. This coincides with the input, as in Hebrew, pairs of strictly identical consonants are frequent root-finally, while pairs of homorganic consonants are rare. Furthermore, knowledge of this difference influences lexical decisions by native listeners of Hebrew (Berent et al., 2004). A similar differentiation between homorganic and identical consonant co-occurrences has also been observed by Frisch and Zawaydeh (2001) with native speakers of Arabic, but they are hesitant to over-interpret the findings as their study was not designed for this comparison. Hence, Hebrew (and possibly Arabic) speakers seem to make a distinction between effects of OCP on identical versus similar consonants.

Another study (Coetzee, 2005) investigated effects of OCP on English listeners. In English, the co-occurrence of identical Cs in /sCVC/ syllables is strongly restricted. However, the constraint only categorically bans co-occurrences of identical labials and identical velars, while there are a few attestations of co-occurrences of identical coronals (e.g., *state*). Coetzee (2005) found that when listening to /sCVC/ syllables, in which one of the Cs was a hybrid sound between two places of articulation, English listeners were generally more likely to perceive this hybrid as non-identical with the second C, and this did not depend on whether the two Cs would have resulted in identical coronals, velars or labials. As stimuli were controlled for lexical statistics, Coetzee suggests that the results are indicative of abstract knowledge of OCP in English listeners (Coetzee, 2005).

This result is in contrast with studies on native listeners of Dutch, whose knowledge of OCP-Place seems to be more detailed, as they use OCP-Labial but not OCP-Coronal for speech processing, as indicated by lexical decisions (Kager and Shatzman, 2007) as well as speech segmentation preferences (Boll-Avetisyan and Kager, 2014, henceforth B&K). In Dutch, co-occurrences of labials but not of coronals are under-attested. Hence, these results suggest a direct link of the use of OCP-Place for processing and input distributions. Moreover, although their study was not designed for making this comparison, B&K report that their participants did not differentiate between homorganic and identical consonant co-occurrences.

(b) *Unsystematic exceptions*: A language might generally restrict the co-occurrence of consonants with shared place, but unsystematically attest (or even over-attest) specific pairs of consonants although they are homorganic. The question of how *unsystematic exceptions* from OCP are dealt with has received little attention in the literature. Alderete et al. (2013) note that among Frisch and Zawaydeh's (2001) stimuli, some consisted of pairs of consonants that are attested in Arabic. Reanalyzing this data, they found that although rating scores were somewhat higher for these items when compared to the rest, the difference was not significant. They are cautious with over-interpreting this result, as Frisch and Zawaydeh's study was not designed to allow a direct comparison. However, as their computational model does not attribute much relevance to exceptions either, they suggest this might be converging evidence that Arabic listeners are not sensitive to phonotactic detail at this level (Frisch and Zawaydeh, 2001, p. 19).

The current study explores how Dutch listeners deal with exceptions from OCP-Place. The Dutch language is an interesting case, because OCP restrictions here are less systematic than those in, for example, the Semitic languages. It could be assumed that effects of abstract OCP are stronger in languages in which effects are more systematic. A question that follows from this possibility is that the knowledge which native listeners of languages such as Dutch use for processing does not necessarily make reference to an abstract representation of OCP, but that they rather refer to detailed representations of specific *n*-phone probabilities. This possibility, which was not addressed by B&K, will be explored in the current study, focusing on both systematic exceptions (i.e., the case of identity versus homorganicity) and unsystematic exceptions (i.e., exceptional over-attestations of specific consonant pairs) from OCP-Place. The following hypotheses can be derived from the literature:

(1) Hypotheses for OCP-Place

(a) *Abstractness hypothesis:*

Knowledge of OCP-Place (or, of the sub-group constraints OCP-Labial, OCP-Coronal and OCP-Dorsal) is abstract; all co-occurrences of consonants with shared [place] are treated equally regardless of their individual degree of attestation.

(b) *Phonotactic detail hypothesis:*

Phonotactic knowledge is detailed; all co-occurrences of consonants with shared [place] are treated differently depending on their individual degree of attestation.

(2) Hypotheses for identity

(a) Abstractness hypothesis:

The identity preference is abstract. All co-occurrences of identical consonants are treated equally even if phonotactic knowledge of non-identical consonant co-occurrences is detailed.

(b) Phonotactic detail hypothesis:

Phonotactic knowledge is detailed; all co-occurrences of identical consonants are treated differently depending on their individual degree of attestation.

Note that (1a) and (1b) as well as (2a) and (2b) need not be mutually exclusive. It is likely that listeners, as suggested in some of the literature reviewed above, rely on both abstract and detailed representations of phonotactics in speech processing. Likewise, the hypotheses regarding (1) OCP-Place and (2). Identity could but must not be mutually exclusive. It could be that knowledge of the identity preference is abstract, while knowledge of OCP-Place is not.

The question raised in the current study is how listeners deal with cases in which cues from abstract versus detailed phonotactic knowledge are in conflict. In order to approach the hypotheses empirically, we will focus on Dutch and OCP-Labial. The paper is structured as follows: First, we will have a detailed look at the distributions of labial consonants in Dutch (Section 2). Second, we will present a new artificial language learning study designed to explicitly test how Dutch listeners deal with (a) specific lexical exceptions from OCP-Labial, and (b) cases of identical consonant co-occurrences in speech processing (Section 3). Third, we will reanalyze the data of B&K to assess whether speech segmentation performance in their study was influenced by knowledge of phonotactic detail rather than abstract OCP-Labial (Section 4).

2. Lexical statistics for co-occurrences of labials in Dutch

The Dutch lexicon has been described to be gradiently restricted by OCP-Place (Boll-Avetisyan and Kager, 2014; Kager and Shatzman, 2007). Both studies used the Observed/Expected (*O*/*E*) ratio (Pierrehumbert, 1993), by which observed counts of a consonant pair in a database is compared to the counts expected if consonants combined at random. The *O*/*E* ratio has been used to illustrate consonant co-occurrence probabilities in many studies (e.g., Adriaans and Kager, 2010; Coetzee and Pater, 2008; Frisch et al., 2004; for a criticism on this measure, see Wilson and Obdeyn, 2009). In B&K's calculations, *O* stood for the number of C_1VC_2 sequences in the corpus (or lexicon). This was divided by *E*, computed as the probability that C_1 occurs in the initial position of CVC, multiplied by the probability that C_2 occurs in the final position of CVC, which is multiplied by the total number of CVC sequence tokens:

$$O/E = N(C_1VC_2)/p(C_1)*p(C_2)*N(CVC)$$

where $N(C_1VC_2)$ is the number of C_1VC_2 sequences, $N(CVC)$ is the number of CVC sequences, $p(C_1)$, $p(C_2)$ are the probability of C_1 and C_2 calculated as proportions of consonants. The values yielded by means of this computation can be interpreted as follows: $O/E > 1$ indicates that a CVC sequence is over-attested, and $O/E < 1$ indicates its under-attestation.

B&K (2014) used the *O*/*E* ratio for calculating the co-occurrence probability of non-adjacent consonants in Dutch in both a corpus of continuous spoken Dutch (Corpus Gesproken Nederlands, henceforth CGN, Goddijn and Binnenpoorte, 2003), and a lexical database of Dutch (CELEX Dutch monomorphemes, Baayen et al., 1995). Their results revealed that Dutch is particularly constrained by OCP-Labial (CELEX: $O/E = 0.45$, CGN: $O/E = 0.58$) and OCP-Dorsal ($O/E = 0.58$, CGN: $O/E = 0.69$), but much less so by OCP-Coronal (CELEX: $O/E = 0.77$, CGN: $O/E = 1.24$), for details see B&K (2014). However, they only provided the reader with the sum of the effects of OCP-Labial, OCP-Coronal and OCP-Dorsal. Information of the *O*/*E* ratios of specific pairs of consonants with shared [place] was not given. The Dutch phoneme inventory includes six labial consonants {/p/, /b/, /m/, /f/, /v/, /w/}, which logically can be combined to 36 labial–labial pairs in CVC sequences (henceforth PVP), such as /pVp/, /pVb/, /pVm/, /bVp/. In order to understand how consistent the effect of OCP-Labial is in the Dutch lexicon, we replicated B&K's calculations.

2.1. Method

We based our calculations on the same 477 CVC types extracted from CELEX Dutch monomorphemes (Baayen et al., 1995) and the same 6312 within- and across word CVC types extracted from CGN as B&K (2014).[1] Furthermore, we calculated the *O*/*E* ratio of 479 within-word CVC types in all word tokens occurring in CGN. These were all CVC sequences independently of whether they occurred word-initially, -medially, -finally or — in CGN — between words. Each occurrence of a C in a database was counted as C_1, no matter whether it was counted as a C_2 before or not. So, from the Dutch phrase *de mannen*/dəmanən/ 'the men', we would extract /dəm/, /man/ and /nən/ from the continuous speech version of CGN. From CELEX and segmented CGN, we would only extract the monomorpheme /man/. Cs that are part of a cluster were also counted, if they were followed by CV. For example, in the word *muts* /myts/ 'hat'), we would count /myt/ (but ignore /s/).

[1] We find an inclusion of CGN continuous speech data interesting, as it may not be that probabilistic phonotactic knowledge is deduced from the lexicon. Even if this is the prevalent position taken in the literature (e.g., Pierrehumbert, 2003), there have also been proposals that phonotactic knowledge as cues for speech segmentation is learned from continuous speech (e.g. Brent and Cartwright, 1996).

2.2. Results and discussion

The results of our calculations indicate that the vast majority of the 36 possible labial–labial pairings in Dutch are indeed under-attested (see Table 1 for monomorphemic word types in CELEX, Table 2 for word tokens in CGN, and Table 3 for continuous speech in CGN). However, five pairings are exempt: the sequences /pVp/, /bVv/ and /bVw/ were found to be over-attested in both, CELEX and CGN. Additionally, /fVb/ is over-attested in CELEX, and /fVf/ in CGN.

In a next step, to understand whether identical consonants co-occur more or less frequently than homorganic labials in Dutch, we averaged the *O*/*E* values of specific pairs of co-occurrences of identical consonants (i.e., /pVp, /bVb/, /fVf/, /vVv/ , /wVw/ and /mVm/). In both the segmented and continuous speech versions of CGN, the *O*/*E* value for pairs of identical labials (CELEX: *O*/*E* = 0.47, CGN words: *O*/*E* = 0.57, CGN continuous: *O*/*E* = 0.68,) was higher than the *O*/*E* value for the sum of labials (CELEX: *O*/*E* = 0.43, CGN words: *O*/*E* = 0.41, CGN continuous: *O*/*E* = 0.50). Note that the small differences in the *O*/*E* ratio between corpora may be due to CELEX being based on written text while CGN is based on spoken text. Another source of the variability might be that calculations in CELEX were based on word types, while calculations in CGN were based on tokens.

To sum up, labial co-occurrence distributions in the Dutch lexicon and in Dutch speech are not systematic: although the majority of consecutive labials are under-attested, a few pairs are over-attested and, hence, exempt from this overall trend. Co-occurrences of identical labials are under-attested. However, especially when considering speech rather than

Table 1
O/*E* values for non-adjacent labial co-occurrences in CVC sequences in Dutch. Calculations are based on CELEX monomorpheme types. Fields shaded in gray indicate over-attestation (i.e., *O*/*E* > 1.0).

C_1	C_2					
	p	b	f	v	w	m
p	1.27	0.69	0.21	0.19	0.39	0.21
b	0.23	0.76	0.18	1.51	3.71	0.53
f	0.14	1.59	0.11	0.15	0.08	0.64
v	0.04	0.21	0.19	0.13	0.17	0.11
w	0.50	0.27	0.55	0.31	0.12	0.34
m	0.20	0.48	0.24	0.11	0.25	0.29

Table 2
O/*E* values for non-adjacent labial co-occurrences in CVC sequences in Dutch; calculations are based on segmented words in the Corpus of Spoken Dutch. Fields shaded in gray indicate over-attestation (i.e., *O*/*E* > 1.0).

C_1	C_2					
	p	b	f	v	w	m
p	1.78	0.87	0.27	0.15	0.15	0.33
b	0.79	0.67	0.63	2.03	1.65	0.29
f	0.06	0.96	1.07	0.51	0.32	0.56
v	0.06	0.80	0.58	0.69	0.27	0.30
w	0.18	0.16	0.20	0.08	0.03	0.37
m	0.07	0.12	0.15	0.20	0.17	0.39

Table 3
O/*E* values for non-adjacent labial co-occurrences in CVC sequences in Dutch; calculations are based on continuous utterances in the Corpus of Spoken Dutch. Fields shaded in gray indicate over-attestation (i.e., *O*/*E* > 1.0).

C_1	C_2					
	p	b	f	v	w	m
p	1.39	0.68	0.46	0.38	0.48	0.52
b	0.57	0.61	0.53	1.26	1.11	0.40
f	0.17	0.68	1.04	0.41	0.31	0.68
v	0.18	0.63	0.51	0.70	0.43	0.49
w	0.24	0.45	0.28	0.21	0.53	0.44
m	0.25	0.39	0.40	0.47	0.58	0.58

Table 4
Predicted segmentation preferences in cases of conflict by hypothesis.

Condition	Over-attestation	Identity	Predicted segmentation outcomes			
			1. OCP-Labial		2. Identity	
			Abstractness hypothesis	Phonotactic detail hypothesis	Abstractness hypothesis	Phonotactic detail hypothesis
bv	Yes	No	/. . .b # v. . ./	/. . . # bv. . ./	–	–
pp	Yes	Yes	/. . .p # p. . ./	/. . . # pp. . ./	/. . . # pp. . ./	/. . . # pp. . ./
bb	No	Yes	/. . .b # b. . ./	/. . .b # b. . ./	/. . . # bb. . ./	/. . .b # b. . ./

the lexicon as a source for deriving distributional knowledge, they are relatively more likely to co-occur than homorganic labials. Previous studies have demonstrated that phonotactic knowledge can be used as a cue for speech segmentation. Generally, listeners are more likely to insert word boundaries between segments if their co-occurrence probability is low than if their co-occurrence probability is high. This effect has been found in studies testing native language phonotactic knowledge of consonant clusters (e.g., McQueen, 1998), vowel harmony (Suomi et al., 1997; Vroomen et al., 1998) and OCP-Labial (B&K, 2014). A question that is still open at this point is whether listeners have a preference for drawing on abstract or detailed phonotactic knowledge when segmenting speech.

The results of the present corpus study indicate that Dutch listeners will encounter PVP sequences in speech in which cues from general OCP-Labial would be in conflict with phonotactic detail and/or constraints on identity. Moreover, they will encounter sequences of identical consonants which would be in conflict with a preference for identity groupings. In light of these results, in principle, all four hypotheses spelled out above regarding knowledge of (1). OCP-Labial as (a) abstract and/or (b) detailed, and (2). Identity as (a) abstract and/or (b) detailed may be justified.

This hypothesized knowledge may possibly lead to the following segmentation strategies by native listeners of Dutch (see Table 4):

If listeners encounter a pair of consonants with shared [place] that is over-attested in the Dutch input (e.g., /bVv/), they would be more likely to insert a word boundary between the two consonants under hypothesis (1a), which states that abstract knowledge of OCP-Labial should be applied irrespective of the degree of attestation of a consonant sequence, while under hypothesis (1b) of phonotactic detail they would not be more likely to insert a boundary.

If listeners encounter a pair of identical consonants that is over-attested in Dutch (e.g., /pVp/), they would not be more likely to insert a boundary under hypothesis (2a) suggesting an abstract identity preference and hypothesis (2b) suggesting knowledge of phonotactic detail, while under hypothesis (1a), an insertion of a boundary would be likely.

If they encounter a pair of identical consonants that is under-attested in Dutch (e.g., /bVb/), then they would not be more likely to insert a boundary under a hypothesis of identity preference (2a), while they should be more likely to insert a boundary under the hypothesis that phonotactic detail (2b) influences segmentation in such situations and/or under the hypothesis (1a) of abstract knowledge of OCP-Labial.

To shed light on this issue, we designed three artificial languages and tested which speech segmentation strategies native listeners of Dutch would apply in the following cases of conflict: (1) /bVv/ (henceforth, bv-condition), an over-attested dependency between homorganic labials conflicting with OCP-Labial, (2) /pVp/ (henceforth, pp-condition), an over-attested dependency between identical labials conflicting with OCP-Labial, and (3) /bVb/ (henceforth, bb-condition), an under-attested dependency between identical labials conflicting with a potential escape hatch from OCP. The predictions that can be stated in light of the hypotheses spelled out above are summarized in Table 4.

3. Experiment

3.1. Method

3.1.1. Participants

For this study, we tested 43 students from the Utrecht University. Recruitment criteria were that they were raised as monolingual native speakers of Dutch, and had no reported hearing or speech deficits. They were compensated for their participation by a fee.

3.1.2. Material

3.1.2.1. Familiarization stimuli. We created three versions of an artificial language in which a CV syllable with a coronal consonant (T) was followed by two syllables starting with labial consonants (P), resulting in a stream of /. . .TPPTPPTPP. . ./. Each language consisted of nine syllables, which were assigned to fixed slots: three P_1 syllables, three P_2 syllables and three

Table 5
Syllable inventory of the artificial languages.

Condition	P_1	P_2	T
bvd-language	/bo/	/va/	/do/
	/bø/	/ve/	/de/
	/bi/	/vy/	/du/
ppt-language	/po/	/pa/	/to/
	/pø/	/pe/	/te/
	/pi/	/py/	/tu/
bbd-language	/bo/	/ba/	/do/
	/bø/	/be/	/de/
	/bi/	/by/	/du/

T syllables (see Table 5). The crucial difference between the languages used here compared to the ones in B&K (2014) is that they had assigned a mix of different consonants to each slot (e.g., P_1 = {/po/, /be/, /ma/}). In the present study, only one specific consonant was assigned to each slot. To test /bVv/ (bv-condition), we designed a *bvd-language*: P_1 syllables started in /b/, P_2 syllables started in /v/ and T syllables started in /d/, resulting in a stream such as /. . .bovadobivydebive. . ./. To test /bVb/ (bb-condition), we designed a *bbd-language*, which was minimally different from the bvd-language: /v/ in slot P_2 was replaced by /b/, resulting in a stream such as /. . .bobadobibydebibe. . ./. To test /pVp/ (pp-condition), we designed a *ppt-language* minimally different from the bbd-language: the voiced consonants were replaced with their voiceless counterparts (i.e., /b/ was replaced by /p/, and /d/ was replaced by /t/), resulting in a stream such as /. . .popatopipytepipe. . ./.

Natural co-occurrence probabilities of the non-adjacent consonant pairs in CELEX and CGN (see Table 6) used in the three languages were as follows: In the bvd-language, /bVv/ is over-attested, while /vVd/ and /dVb/ are under-attested. In the ppt-language, /pVp/ is over-attested, while /pVt/ and /tVp/ are under-attested. In the bbd-language, /bVb/ is under-attested, while /bVd/ and /dVb/ are over-attested.

In order to reduce the possibility that listeners would use probabilistic information other than the co-occurrence probability of the non-adjacent consonants as a cue for segmentation, we based the syllable selection and the way syllables were combined on controls for lexical statistics. As said, syllables in the third slot started with T syllables. As Dutch only has two velars (= K, namely /x/ and /k/) that are licit in syllable onsets, it was easier for us to control for lexical statistics by using T rather than K. First, we controlled for positional syllable frequencies. That is, we calculated the frequency with which PV and TV syllables would occur word-initially, word-medially or word-finally in Dutch words listed in CELEX, and selected those that were most evenly balanced to avoid positional syllable frequencies to give a cue to segmentation (see Supplement 1). Furthermore, we controlled for transitional probabilities between phonemes. That is, the probability by which a specific C was followed by a specific V and a specific V was followed by a specific C was balanced between PPT and PTP test items in each condition (see Supplement 2). Next, we controlled all possible three-syllabic chunks (i.e., PPT, PTP, TPP parsings) for lexical neighborhood density by applying Luce's (1986) criterion, according to which a lexical neighbor is any word that differs from a given item by a single segmental change, for example by deletion, addition, permutation or alteration. The syllables were selected and combined such that no three-syllabic string had a lexical neighbor in CELEX. Lastly, we aimed at controlling for cohort density (CD), which turned out to be difficult (see Supplement 2). CD was defined as the sum of the logged frequencies of the set of words in CELEX starting with the same three phonemes. Among our conditions, CD was only truly controlled for in the ppt-language. In the bbd-language, #bd had on average slightly higher CD than #bb (which works against the *escape hatch for identity hypothesis*), and in the bvd-language, #bv had slightly higher CD than #vd (which works against the *abstractness hypothesis*). Each artificial language stream consisted of 2700 syllables. Hence, the stream can be segmented into 900 PTP, 900 PPT or 900 TPP tri-syllabic tokens. In each of the three artificial languages, the syllable order was pseudo-random with the only constraints being that syllables had to fall into their fixed slot, and that no tri-syllabic string was repeated in immediate

Table 6
Natural *O*/*E* values of the non-adjacent consonants used in the artificial languages; calculations based on monomorphemic word types in CELEX, word tokens in CGN and continuous speech in CGN.

Database	pp	pt	tp	bb	bd	db	bv	vd	dv
CELEX	1.27	1.00	0.71	0.76	1.61	0.96	1.51	0.59	0.68
CGN words	1.78	0.62	0.81	0.67	1.32	0.31	2.03	0.63	0.93
CGN cont	1.39	0.83	1.10	0.61	1.32	1.21	1.26	0.49	0.24

succession. The languages were synthesized with a male Dutch voice using the ''nl2'' diphone database (provided by FLUENCY) and the MBROLA speech synthesizer (Dutoit et al., 1996). Syllables had a duration of 232 ms and pitch was monotonous over the whole stream ($F0$ = 100 Hz). At the beginning of the familiarization stream, intensity faded in for the first 5 s, and the same fade-off was used at the end of the familiarization stream in order to prevent the endpoints from giving a cue for segmentation. When played, the stream was 10 min long.

In order to reduce the possibility that listeners would use probabilistic information other than the co-occurrence probability of the non-adjacent consonants as a cue for segmentation, we based the syllable selection and the way syllables were combined on controls for lexical statistics. As said, syllables in the thirds slot started with T syllables. As Dutch only has two velars (= K, namely /x/ and /k/) that are licit in syllable onsets, it was easier for us to control for lexical statistics by using T rather than K. First, we controlled for positional syllable frequencies. That is, we calculated the frequency with which PV and TV syllables would occur word-initially, word-medially or word-finally in Dutch words listed in CELEX, and selected those that were most evenly balanced to avoid positional syllable frequencies to give a cue to segmentation (see Supplement 1). Furthermore, we controlled for transitional probabilities between phonemes. That is, the probability by which a specific C was followed by a specific V and a specific V was followed by a specific C was balanced between PPT and PTP test items in each condition (see Supplement 2). Next, we controlled all possible three-syllabic chunks (i.e., PPT, PTP, TPP parsings) for lexical neighborhood density by applying Luce's (1986) criterion, according to which a lexical neighbor is any word that differs from a given item by a single segmental change, for example by deletion, addition, permutation or alteration. The syllables were selected and combined such that no three-syllabic string had a lexical neighbor in CELEX. Lastly, we aimed at controlling for cohort density (CD), which turned out to be difficult (see Supplement 2). CD was defined as the sum of the logged frequencies of the set of words in CELEX starting with the same three phonemes. Among our conditions, CD was only truly controlled for in the ppt-language. In the bbd-language, #bd had on average slightly higher CD than #bb (which works against the *escape hatch for identity hypothesis*), and in the bvd-language, #bv had slightly higher CD than #vd (which works against the *abstractness hypothesis*). Each artificial language stream consisted of 2700 syllables. Hence, the stream can be segmented into 900 PTP, 900 PPT or 900 TPP tri-syllabic tokens. In each of the three artificial languages, the syllable order was pseudo-random with the only constraints being that syllables had to fall into their fixed slot, and that no tri-syllabic string was repeated in immediate succession. The languages were synthesized with a male Dutch voice using the ''nl2'' diphone database (provided by FLUENCY) and the MBROLA speech synthesizer (Dutoit et al., 1996). Syllables had a duration of 232 ms and pitch was monotonous over the whole stream ($F0$ = 100 Hz). At the beginning of the familiarization stream, intensity faded in for the first 5 s, and the same fade-off was used at the end of the familiarization stream in order to prevent the endpoints from giving a cue for segmentation. When played, the stream was 10 min long.

3.1.2.2. Test stimuli. For the test phase, PPT and PTP items were synthesized for each condition. Within each condition, 27 PTP and 27 PPT sequences were used as test items. We used every PTP or PPT sequence that occurred in the artificial languages as a test word. In the bv-condition, this contrasted vdb (e.g., /vydobo/) and bvd (e.g., /bovede/) items, the pp-condition contrasted ptp (e.g., /pytopo/) and ppt (e.g., /popete/) items, and the bb-condition contrasted bdb (e.g., /bydobo/) and bbd (e.g., /bobede/) items. All stimuli are listed in Supplement 2. TPP items were not included.[2] Each test item occurred twice within a test phase, once as a first and once as a second item in a pair, resulting in 54 trials per condition. The test pairs were combined at random with the prerequisite that two items only made a pair once, and exactly half of the test pairs started with one type of stimulus.

3.1.3. Procedure

The procedure, originally borrowed from Peña et al. (2002), was paralleled with B&K (2014): Participants were tested individually in a sound-attenuated booth. They listened to the stimuli through headphones. Prior to entering the familiarization phase, participants had to carry out a little pre-test in which they had to identify a CV syllable. This pre-test had the purpose of acquainting participants with the synthetic speech and with the two-alternatives-forced-choice task of the test phase. After this pre-test, participants were instructed on screen that they would listen to an artificial language for the next 10 min and were asked to listen carefully, because they would later be tested on their knowledge about the words of this language. After the familiarization phase, they were tested in a two-alternative forced choice task, in which they heard a PTP and a PPT item. They had to indicate via a mouse-click on a button on the screen whether the first or the

[2] B&K (2014) had three conditions: PTP-PPT, PTP-TPP and PPT-TPP. In the first two conditions, listeners showed a PTP-preference, and in the last condition, they had no significant preference. This suggests that listeners do not show a preference for items unless they have segmented them from the artificial language. In the current study, we started out by testing PTP-PPT, so both item types start with labials, reducing the possibility that judgments would be driven by item onsets. An addition of control conditions including TPP may, following the results of our prior study, only become relevant if participants do not display a preference for one over the other segmentation in a test phase contrasting PTP-PPT.

second of the two items was more likely to be a word from the language they just heard. Test items within a pair were presented with an inter-stimulus interval of 500 ms.

Using a between-subjects-design, participants were assigned to one out of three conditions (ppt-language: $N = 15$; bbd-language: $N = 14$; bvd-language: $N = 14$). Within a condition, all participants received the same list of test pairs. The order of the presentation of the test pairs was randomized.

3.1.4. Data processing and analysis

One participant who was tested in the bb-condition was excluded from the analysis, because it turned out after test she had not been raised monolingually. Two generalized linear mixed models with logits (Jaeger, 2008) were fitted using the statistics program ''R'' version 3.0.2. (R Core Team, 2012) with the packages lme4 (Bates et al., 2012) for the mixed model analyses and ggplot2 (Wickham, 2009) for the figures.

In both models, the dependent variable was response type (0 = PPT, 1 = PTP), condition (bvd-, ppt-, or bbd-language) was a fixed factor, and participants and items were specified as random factors. The first model was to assess whether segmentation preferences in each condition differed from chance. Here, negative coefficients (= β) indicate performance below chance (i.e., a preference for a PPT-segmentation), and positive coefficients indicate performance above chance (i. e., a PTP-segmentation). In the second model we compared segmentation preferences between conditions. For this, a successive difference contrast (using the contr.sdif() function in ''R'') was coded. This is an orthogonal contrast which uses the grand mean in the intercept, and compares levels in a step-wise fashion (1 is compared to 2, and 2 is compared to 3); coefficients indicate the difference between levels. In our analysis, segmentation preferences in the bvd-language were compared to the ppt-language (hence, comparing the two languages with over-attested PP), and the ppt-language was compared to the bbd-language (hence, comparing the two languages including identical Ps). Furthermore, as the results could potentially be confounded by cohort densities (CD), we added this as a covariate. As items were presented in pairs of PPT versus PTP items, we subtracted the density values of PVP cohorts from those of PVT cohorts. Hence, a negative CD value reflects that in a given test pair, the PPT item had a higher CD than the PTP item. The CD variable was centered around its mean to reduce collinearity. As model comparisons revealed that the insignificant interaction of CD * Condition did not account for additional variance, the interaction term was not included in the final model.

3.2. Results and discussion

3.2.1. Results

Fig. 1 illustrates that participants were more likely to prefer bvd (60%) than vdb items in the bv-condition, and ppt (62%) than ptp items in the pp-condition. In the bb-condition, however, they were less likely to prefer bdb (40%) than bbd items. Results of the tests against chance revealed that preferences for a PPT over PTP segmentation was significant in both the bv- and pp-condition, while the preference for a PTP over PPT segmentation in the bb-condition was only marginally significant (see Table 7 for the statistics).

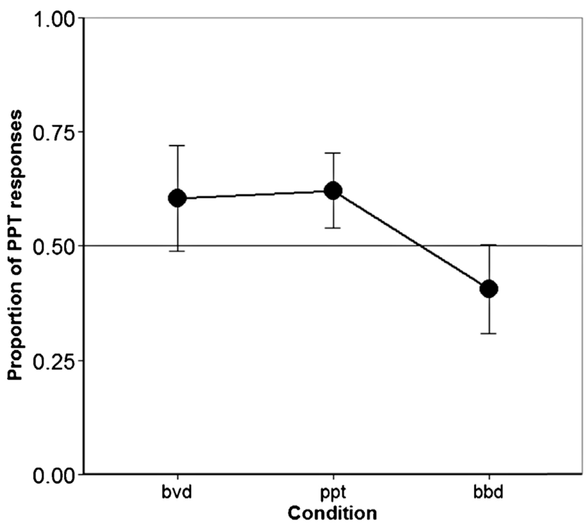

Fig. 1. Proportion of PPT responses ±two standard errors (indicated by the error bars) by condition (bvd-language, ppt-language, and bbd-language).

Table 7
Tests of response preferences for each condition against chance. Coefficients (*β*) and their standard errors (*SE*) are logit transformations.

Estimates of fixed effects

Parameters	*β*	*SE*	*z*	*p*
bvd-language	0.54	0.28	1.99	< 0.05
ppt-language	0.60	0.26	2.30	< 0.05
bbd-language	−0.50	0.28	−1.78	0.08, n.s.

Estimates of random effects

Groups	Name	Variance	*SD*
Items	Intercept	0.019	0.138
Participants	Intercept	0.914	0.956

Table 8
Results of the comparisons between conditions. Coefficients (*β*) and their standard errors (*SE*) are logit transformations.

Estimate of fixed effects

Parameters	*β*	*SE*	*z*	*p*
Intercept	0.22	0.17	1.34	0.19, n.s.
ppt-language–bvd-language	0.14	0.39	0.36	0.72, n.s.
bbd-language–ppt-language	−1.09	0.39	−2.80	< 0.01
Cohort density	−0.0035	0.0014	−2.43	< 0.05

Estimates of random effects

Groups	Name	Variance	*SD*
Items	Intercept	0.12	0.35
Participants	Intercept	0.96	0.98

Comparisons between conditions in a second model revealed a significant difference between the pp- and the bb-condition: there were significantly more PPT responses at test in the case of the ppt-language than in the case of the bbd-language. However, there was no significant difference between the bv- and the pp-condition. Furthermore, there was a main effect of cohort density: participants were more inclined to select a test item if it had a higher cohort density across conditions (see Table 8 for the statistics).

3.2.2. Discussion

That cohort density has an effect on speech segmentation has been known for quite some time (e.g., Norris et al., 1995). However, the results we found beyond the effect of CD, that is, the effects of condition, can inform us about how OCP-Place is represented in the minds of speakers of a language that gradiently restricts its lexicon by means of this constraint. The material in this experiment was constructed such that it would allow us to assess four hypotheses that were spelled out in the introduction. Below, we will discuss to what extent the results are or are not in line with the predictions that were derived from the possible hypotheses (see Table 9 for an overview).

Table 9
Predicted segmentation outcomes by hypothesis per condition. Segmentation outcomes are in bold. Hypotheses rejected by the segmentation outcomes are shaded in gray.

Condition	Predictions for possible segmentation outcomes			
	1. OCP-Labial		2. Identity	
	(a) Abstractness hypothesis	(b) Phonotactic detail hypothesis	(a) Abstractness hypothesis	(b) Phonotactic detail hypothesis
bvd-language	/. . .b # vdb # vd. . ./	**/. . . bvd # bvd. . . /**	–	
ppt-language	/. . .p # ptp #pt. . ./	**/. . . ppt # ppt. . . /**	**/. . . ppt # ppt. . . /**	**/. . . # pp. . . /**
bbd-language	**/. . . b # bdb # bd. . . /**	**/. . . b # bdb # bd. . . /**	/. . .bbd # bbd. . ./	**/. . . b # b. . . /**

The *abstractness hypothesis* of OCP-Labial (1a) predicts that listeners would be more likely to insert word boundaries between consecutive labials independent of whether they are exceptions from OCP-Labial or not. Hence, listeners should have been more likely to segment /. . .b # vdb # vd. . ./ in the bv-condition, /. . .p # ptp # pt. . ./ in the pp-condition, and /. . .b # bdb # bd. . ./ in the bb-condition. However, instead, they had a segmentation preference for /. . .bvd # bvd. . ./ in the bv-condition, for /. . .ppt # ppt. . ./ in the pp-condition, and in the bb-condition, they only had a marginal preference for /. . .b # bdb # bd. . ./. Hence, in none of the conditions were segmentation results as predicted by the *abstractness hypothesis* of OCP-Labial. That is, the current results do not support a use of abstract knowledge of OCP-Labial in speech segmentation by Dutch listeners. Instead, the results of this experiment were, across conditions, in line with the predictions derived from the *phonotactic detail hypothesis*, by which listeners would only be more likely to insert word boundaries between consecutive labials if they are under-attested in the native language's lexicon, but not if they are over-attested. And indeed, listeners preferred /. . .bvd # bvd. . ./ in the bv-condition, /. . .ppt # ppt. . ./ in the pp-condition, and /. . .b # bdb # bd. . ./ in the bb-condition.

We also explored the hypothetical effects of an *escape hatch for identity* according to which listeners might treat pairs of identical segments differently from pairs of homorganic consonants (hypothesis (2a)). As pairs of identical consonants seem to be less under the influence of OCP-Labial than pairs of non-identical consonants, it would be predicted that listeners would not be likely to insert boundaries between identical labials (2a). This hypothesis may seem to be confirmed when considering the results of the tests against chance, which suggest that listeners have segmented /. . .ppt # ppt. . ./ in the pp-condition, and had no clear segmentation preference in the case of the bbd-language in the bb-condition. However, the significant difference between segmentation preferences in the bbd- versus the ppt-language suggests they had a stronger preference for keeping over-attested /pVp/ than under-attested /bVb/ intact. This, again, supports the *phonotactic detail hypothesis* (2b). Even if the possibility remains that phonotactic detail affected segmentation in conjunction with an escape hatch from OCP, the results do not suggest that the Dutch listeners have applied a *general escape hatch from OCP-Labial* strategy when segmenting the two languages.

These results suggest that in this speech segmentation task, Dutch listeners preferred attending to phonotactic detail rather than to abstract phonotactic constraints if cues from abstract versus detailed phonotactics are in conflict. A question that arises from these results is whether Dutch listeners who participated in B&K's (2014) study were also guided by phonotactic detail. Posing this question is justified under the hypothesis that listeners have access to phonotactic detail and abstract constraints. This possibility will be assessed in Section 4.

4. Reanalyzing Boll-Avetisyan and Kager's (2014) data

B&K (2014) familiarized Dutch adults with an artificial speech stream in which two CV syllables with a labial consonant (P) were followed by a CV syllable with a coronal consonant (T, e.g., /. . .pamatumomatubibetu. . ./). Transitional probability between syllables in this stream was always 0.33. The material was highly controlled for lexical statistics to reduce effects from cues other than OCP-Labial on speech segmentation. In a consecutive forced choice task, participants more readily recognized PTP items than PPT or TPP items as part of the artificial language.

B&K's artificial languages included nine different combinations of PVP, i.e., pVp, pVb, pVm, bVp, bVb, bVm, mVp, mVb, mVm. Hence, besides or in addition to abstract OCP-Labial, their segmentation preferences could also have been influenced by (a) an escape hatch for identity or by (b) phonotactic detail. B&K considered the possibility that sequences of identical consonant co-occurrences (i.e., pVp, bVb, and mVm) may have been treated differently from homorganic consonants in their study. Their mixed model analyses, however, did not suggest that this factor, which did not reach significance, was relevant for accounting for the variance in speech segmentation preferences. However, they did not consider whether knowledge of phonotactic detail could also explain the results. If that were the case it would be predicted that listeners would more often select PPT and TPP test items as being part of the artificial language if they contained more likely PVP sequences (e.g., test items *popena* and *dopapo*) than if they did not (e.g., test items *mobene* and *tupomi*). To clarify this, we reanalyzed B&K's data.

4.1. Data processing and analysis

For our reanalysis, we used the data of Condition 1 and Condition 2 in B&K's (2014) Experiment 1 ($N = 71$). In Condition 1, PTP and PPT items were contrasted in the test phase, and in Condition 2, PTP and TPP items were contrasted. The dependent variable was response (1 = PTP response, 0 = PPT/TPP response). Fixed factors were Condition (1 versus 2) and a continuous predictor reflecting the probability of PVP using PVP's *O*/*E* values as measured in CELEX. For the fixed factor Condition, we coded a successive difference contrast. This orthogonal contrast adds the grand mean to the intercept, and the coefficient (β) indicates the difference score between the two conditions. The continuous predictor was centered around its mean to reduce collinearity.

Table 10
Results of the mixed model reanalyzing B&K's (2014) data adding the likelihood of PVP occurrences (in CELEX) as a covariate. Coefficients (β) and their standard errors (*SE*) are logit transformations.

Estimate of fixed effects

Parameters	β	*SE*	*z*	*p*
Intercept	0.23	0.08	2.81	< 0.01
Condition	−0.03	0.16	−0.16	0.87, n.s.
O/*E*s of PVP	0.13	0.17	0.76	0.45, n.s.
Condition * *O*/*E*s of PVP	−0.01	0.34	−0.03	0.97, n.s.

Estimates of random effects

Groups	Name	Variance	*SD*
Items	Intercept	0.34	0.58
Participants	Intercept	0.21	0.45

4.2. Results and discussion

As a result (see Table 10), only the intercept was significant. The intercept's positive coefficient indicates that overall participants preferred PTP items over their counterparts, that is, PPT items in Condition 1, and TPP items in Condition 2. Furthermore, there was no main effect of condition, that is, the preference for PTP items was about the same independent of whether they were contrasted with PPT or TPP items. The preference for PTP items at test in both the conditions indicates that they have inserted boundaries between consecutive labials during segmentation (e.g., /. . .pa-matumo-matubi-betu. . ./), which is indicative of their use of OCP-Labial as a speech segmentation cue. To this part, we replicate what we knew from B&K. In addition to this, the results of the new model including the factor ''*O*/*E*s of PVP'' showed that there was neither a main effect of the likelihood of PVP nor an interaction of condition and the likelihood of PVP. Furthermore, a model comparison between the presented model and a second model excluding the factor ''*O*/*E*s of PVP'' showed that an inclusion of this predictor does not significantly improve the model fit. That is, it does not catch any unaccounted variance in the data (χ^2 ($df = 2$) = 0.61, $p = 0.74$).

In light of the results of the new experiment presented in the current study, it may come as a surprise that the co-occurrence probabilities of PVP did not predict the segmentation preferences in our previous data. Note that B&K's material was otherwise strongly controlled for lexical statistics. Although we have to be careful not to put too much weight on non-significant outcomes, potentially, this may suggest that the participants in B&K's experiment were really guided by abstract knowledge of OCP-Labial. The seemingly contradictory outcomes of B&K's versus our experiment make sense when considering the difference in task demands of these experiments. Potential influences of task demands and also the potential role of the native language of the participants on the use of OCP-Place on speech segmentation will be discussed in Section 5.

5. General discussion

In this paper, we aimed to shed light on the question of how gradient phonotactic knowledge influences speech processing. More specifically, we investigated whether Dutch listeners have a preference for drawing on abstract or detailed phonotactic knowledge as a cue for speech segmentation, focusing on non-adjacent labial co-occurrences. Our calculation of lexical statistics over a continuous speech corpus and a lexical database of Dutch indicated effects of OCP-Labial generally hold, yet are not totally systematic: First, five specific pairings occur more often than expected in the Dutch lexicon in spite of a general statistical effect of OCP-Labial. Second, on average, identical labials are more likely to co-occur than homorganic labials.

We then set out to explore whether phonotactic detail or abstract OCP-Labial influences segmentation, and whether co-occurrences of identical labials would give different speech segmentation cues than co-occurrences of homorganic labials. Across conditions, the results of the experiment suggest that participants have relied on knowledge of phonotactic detail when segmenting artificial languages. These results raised our interest in the source of an effect in B&K's (2014) study, which found evidence for an influence of OCP-Labial on speech segmentation by Dutch listeners, but had not considered that segmentation preferences might have been guided by knowledge of phonotactic detail. When reanalyzing B&K's data, however, we did not find that their segmentation results were influenced by phonotactic detail.

The present experiment was designed to test whether abstract or detailed phonotactic knowledge influences segmentation when cues from both collide. Results suggest that phonotactic detail is more relevant than abstract knowledge. However, the reanalysis of B&K's data casts this conclusion into doubt. In the following discussion, we will first consider that these studies differed with respect to a number of factors with potential impact on processing demands: (1) they used different materials, (2) they used different experimental tasks, and (3) they used participants with different native languages. After this, we will discuss the limitations of the current study and will propose some directions for future research.

5.1. Effects of stimulus material

A number of studies have suggested that listeners can use both lower-level acoustic/phonetic and higher-level abstract phonological knowledge for speech processing. Differences in processing strategies, that is, differences in whether listeners rely more on acoustic/phonetic or on abstract phonological knowledge, can be invoked by differences in task demands. For example, task demands can be modulated by stimulus complexity: the more complex the material that needs to be processed, the higher the processing load. Such task dependent effects on processing strategies have, for example, been found in studies on prosodic processing, in which listeners were found to rely on more abstract representations when stimulus complexity was increased by an increased token variability or speaker variability (e.g., Dupoux et al., 1997, 2001b) or segment variability (e.g., Bhatara et al., 2013, 2015).

Having this in mind when considering material differences in the experiments between the current and prior studies, we observe the following: In the current study, the stimulus material was rather simple. For more than 10 min, listeners are continuously exposed to a set of only two (or three) different consonants and six vowels which are presented in the same order. While we find evidence for the use of phonotactic detail in the current study, all other studies that suggest influences of abstract OCP have used material that is relatively more complex. The material in B&K's study included a mix of different consonants and vowels presented in a larger variety of possible combinations. Moreover, both Berent and colleagues' (1997, 2001b, 2002) as well as Frisch and Zawaydeh's (2001) stimulus material included an even higher number of different phonemes in various positions. Furthermore, Coetzee's (2005) material, although overall relatively simple (e.g., when compared to the material used in wellformedness judgment tasks by Frisch and Zawaydeh as well as Berent and colleagues), also used a variety of different consonant pairs including labials, coronals and dorsals in different syllable positions in a within-subjects design. Hence, this might suggest that abstract constraints affect processing when the speech input is of a certain degree of complexity and higher processing demands invoke a reliance on abstract representations. When, however, processing demands are relatively simple, then effects of phonotactically detailed lower-level representations can be invoked. That is to say that Dutch listeners are likely to have both, abstract knowledge of OCP-Labial and phonotactically detailed knowledge of the specific co-occurrence probabilities of labials, and can be influenced by both types of representations depending on processing demands.

5.2. Task effects

It is well-known that differences between experimental tasks can lead to differences in the results, as they may invoke different demands on processing. To give a simple example, Coetzee (2009) showed that wellformedness judgment tasks can elicit decisions reflecting categorical phonotactic knowledge if the task is to rate the wellformedness of one nonword at a time, while they can elicit decisions reflecting gradient phonotactic knowledge when the task is to compare two nonwords at a time.

When comparing the task used in our study to previous studies, differences in processing demands become evident. Most studies that suggest a psychological reality of abstract phonological knowledge used wellformedness judgments (Berent and Shimron, 1997; Coetzee, 2009) and lexical decision tasks (e.g., Berent et al., 2001b; Shatzman and Kager, 2007). Of course it has often been shown that decisions in these tasks are also affected by lower-level probabilistic information and lexical neighborhood density (e.g., Bailey and Hahn, 2001; Vitevitch and Luce, 1998, 1999). Nonetheless, it might be that — in contrast to a speech segmentation task as employed in the current study — wellformedness judgment and lexical decision tasks directly invite listeners to make comparisons between a nonword and the whole lexicon, which might render a processing strategy relying on abstract representations as relatively more successful. It might even be that this comparison between a nonword and the full lexicon only gives the impression of an influence of abstract representations, while in reality, they spontaneously draw analogies, even if nonwords are controlled for lexical statistics. Some evidence for this comes from a study by Frisch et al. (2001): When reanalyzing Frisch and Zawaydeh's (2001) data, they found that specific nonwords with a highly similar neighbor were judged as more wellformed than nonwords with the same distributional properties that did not have a highly similar neighbor.

In our study, however, listeners were required to listen to a speech stream of recurrent syllable strings. The repetition of the syllables may have led them to pay attention to specific phoneme transitions, and they probably have tried to memorize specific syllable combinations. It seems reasonable that such a task invokes an activation of phonotactically detailed knowledge.

5.3. Effects of language background

The OCP restricts the lexicons of the worlds' languages to different degrees. In some languages, OCP-effects are largely systematic. Examples are Hebrew and Arabic, where the ban of identical and homorganic consonant co-occurrences is very regular. In other languages, the ban of identical and/or homorganic consonant co-occurrences is less systematic. Dutch and English are examples. English strongly restricts co-occurrences of identical consonants in #/sCVC/; however, restrictions on other pairs of homorganic consonants are rather probabilistic. Similarly, Dutch only displays a trend toward under-attesting homorganic and identical consonant co-occurrences, which again depends on [place] with stronger restrictions against co-occurring labials and dorsals than against co-occurring coronals. Other languages such as Mandarin Chinese, Guarani, Wantoat and Tatar do not even display any effects of OCP-Place (Boll-Avetisyan, 2012b; Graff, 2012).

It is important to consider that a use of the OCP for processing should only be justified if this potentially facilitates lexical retrieval. Hence, starting with one extreme case, native speakers of languages which are not restricted by OCP-Place should not use this constraint for speech processing. This has, indeed, been found: native listeners of Mandarin Chinese do not use OCP-Labial for speech segmentation unless they are advanced second language learners of Dutch (Boll-Avetisyan, 2012b). Now, considering the other extreme, namely languages which categorically exclude words with pairings of homorganic consonants (even if they allow for a few exceptions), in such cases, native speakers of such languages should experience facilitation by the possibility to rely on an abstract representation of the OCP. For languages like Dutch, which do not systematically under-represent co-occurrences of identical or homorganic consonants, a reliance on abstract OCP only would not be ideal.

5.4. Abstract constraints versus phonotactic detail

There is much evidence for both distributional and abstract phonotactic knowledge in the literature. This suggests that hybrid models adding a stochastic component to phonological grammar might offer the best reflection of the representations in listeners' minds. The issue of constraint specificity versus constraint generality has arisen in much previous work (e.g., Adriaans and Kager, 2010; Albright and Hayes, 2003; Hayes, 1999; Hayes and Wilson, 2008). One of these studies dealt with the question of whether abstract constraints or detailed phonotactic detail give cues to speech segmentation: Adriaans and Kager (2010) computationally simulated the acquisition of phonotactic speech segmentation cues. As a result, maximal speech segmentation performance was yielded when their model was trained to generalize over phonetic features, but at the same time also learned the probability of specific phoneme combinations, and this phonotactic detail helped the model to set correct boundaries in cases in which a specific combination of phonemes was exempt from the abstract constraint.

The results of the present experiment only provide evidence for Dutch listeners to rely on phonotactic detail in speech segmentation. That is, in our speech segmentation task, Dutch listeners preferred attending to cues from phonotactic detail when these were in conflict with cues from abstract OCP-Labial or an abstract identity preference. However, as indicated in our discussion above, these results may have to be attributed to the processing demands posed by our task and stimulus material, and they may not be generalizable to speakers of languages other than Dutch. Hence, we believe that abstract representations of phonotactic knowledge could help listeners hold linguistic input in short-term memory during processing, while phonotactic detail could help resolve ambiguous cues.

If future studies want to investigate this issue further, they should consider the following directions: First, they should assess whether an effect of phonotactic detail was partly suppressed by the presence of knowledge of abstract OCP-Labial or an abstract identity preference. To assess this, it will be necessary to construct stimulus material that allows for the following comparisons:

(1) For assessing whether abstract knowledge of OCP-Labial influences segmentation beyond detail, it will be necessary to contrast sequences of homorganic non-homorganic consonants of comparable co-occurrence probability. Ideally this should be done for both over- and under-attested pairs of homo- and heterorganic consonants. If there is an effect of abstract OCP-Labial, then listeners should be more likely to insert boundaries between labials more often if they are homorganic than if they are heterorganic.

(2) For assessing whether an abstract identity preference influences processing beyond detail, it will be necessary to contrast sequences of identical with homorganic, non-identical consonants of comparable co-occurrence probability. If there is an effect of an abstract identity preference, then listeners should be more likely to insert boundaries between labials more often if they are non-identical than if they are identical.

A second direction future studies should take is to directly manipulate factors that might lead to differences in processing demands, such as the complexity of the material. Moreover, it would be desired if future studies involved cross-linguistic comparisons in order to better understand to which degree a reliance on abstract versus detailed phonotactic knowledge in speech processing depends on the phonotactic system as a whole.

To sum up, even if the current study only provides direct evidence for a use of phonotactic detail as a cue for speech segmentation, altogether, there are good reasons to believe that listeners might have access to a combination of both abstract OCP-Labial and phonotactic detail. It would be desirable if future studies further addressed this issue.

Acknowledgements

We would like to thank Frans Adriaans, Adamantios Gafos, Tom Lentz and the audiences at *Manchester Phonology Meeting 2008*, *Laboratory Phonology Conference 11, 2008*, *Workshop on Phonological Similarity, NELS 40, 2010*, *2nd International Workshop on Phonotactics and Phonotactic Modeling, 2013*, and *Conference on Empirical Approaches to the Phonological Structure of Words, 2014*, for discussions and feedback. Furthermore, we thank Tanner Sorensen for proofreading the manuscript. This work was supported by the Netherlands Organisation for Scientific Research Grant No. 277-70-001 awarded to RK.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.lingua.2015.11.008.

References

Adriaans, F., Kager, R., 2010. Adding generalization to statistical learning: The induction of phonotactics from continuous speech. J. Mem. Lang. 62, 311–331.
Albright, A., Hayes, B., 2003. Rules vs. analogy in English past tenses: a computational/experimental study. Cognition 90, 119–161.
Alderete, J., Tupper, P., Frisch, S.A., 2013. Phonological constraint induction in a connectionist network: learning OCP-Place constraints from data. Lang. Sci. 37, 52–69.
Baayen, R.H., Piepenbrock, R., Gulikers, L., 1995. The CELEX Lexical Database (CD-ROM). Linguistic Data Consortium, University of Pennsylvania, Philadelphia.
Bailey, T.M., Hahn, U., 2001. Determinants of wordlikeness: phonotactics or lexical neighbourhoods? J. Mem. Lang. 44, 568–591.
Bates, D., Maechler, M., Bolker, B., 2012. lme4: Linear Mixed-Effects Models Using S4 Classes. R Package Version 0.999999-0.
Berent, I., Shimron, J., 1997. The representation of Hebrew words: evidence from the Obligatory Contour Principle. Cognition 64, 39–72.
Berent, I., Shimron, J., 2003. Co-occurrence restrictions on identical consonants in the Hebrew lexicon: are they due to similarity? J. Linguist. 39, 31–55.
Berent, I., Everett, D.L., Shimron, J., 2001a. Do phonological representations specify variables? Evidence from the Obligatory Contour Principle. Cognit. Psychol. 42, 1–60.
Berent, I., Shimron, J., Vaknin, V., 2001b. Phonological constraints on reading: evidence from the Obligatory Contour Principle. J. Mem. Lang. 44, 644–665.
Berent, I., Marcus, G.F., Shimron, J., Gafos, A.I., 2002. The scope of linguistic generalizations: evidence from Hebrew word formation. Cognition 83, 113–139.
Berent, I., Vaknin, V., Shimron, J., 2004. Does a theory of language need a grammar? Evidence from Hebrew root structure. Brain Lang. 90, 170–182.
Berent, I., Wilson, C., Marcus, G., Bemis, D.K., 2012. On the role of variables in phonology: remarks on Hayes and Wilson. Linguist. Inq. 43, 97–119.
Berkley, D., 1994. The OCP and gradient data. Stud. Linguist. Sci. 24, 59–72.
Bhatara, A., Boll-Avetisyan, N., Unger, A., Nazzi, T., Höhle, B., 2013. Native language and stimulus complexity affect rhythmic grouping of speech. J. Acoust. Soc. Am. 134 (5), 3828–3843.
Bhatara, A., Boll-Avetisyan, N., Agus, T., Höhle, B., Nazzi, T., 2015. Language experience affects grouping of musical instrument sounds. Cognitive Sci.. http://dx.doi.org/10.1111/cogs.12300in press.
Boersma, P., Hayes, B., 2001. Empirical tests of the gradual learning algorithm. Linguist. Inq. 32 (1), 45–86.
Boll-Avetisyan, N., 2012a. Probabilistic phonotactics in lexical acquisition: The role of syllable complexity. In: Hoole, P., Bombien, L., Pouplier, M., Mooshammer, C., Kühnert, B. (Eds.), Consonant Clusters and Structural Complexity. Mouton de Gruyter, Interface Explorations Series, Berlin/Boston.
Boll-Avetisyan, N., 2012b. Phonotactics and its Acquisition, Representation, and Use: An Experimental-Phonological Study. LOT Dissertation series 298. LOT Publications, Utrecht.
Boll-Avetisyan, N., Kager, R., 2014. OCP-Place in speech segmentation. Lang. Speech 57 (3), 394–421.
Brent, M.R., Cartwright, T.A., 1996. Distributional regularity and phonotactic constraints are useful for segmentation. Cognition 61, 93–125.
Chomsky, N., Halle, M., 1968. The Sound Pattern of English. Harper & Row, New York.
Coetzee, A., 2005. The OCP in the perception of English. In: Frota, S., Vigario, M., Freitas, M.J. (Eds.), Prosodies. Mouton de Gruyter, New York, pp. 223–245.
Coetzee, A., 2009. Grammar is both categorical and gradient. In: Parker, S. (Ed.), Phonological Argumentation. Equinox, London.

Coetzee, A., Pater, J., 2008. Weighted constraints and gradient restrictions on place co-occurrence in Muna and Arabic. Nat. Lang. Linguist. Theory 26, 289–337.
Coleman, J., Pierrehumbert, J.B., 1998. Stochastic phonological grammars and acceptability. In: Proceedings of the 3rd Meeting of the ACL Special Interest Group in Computational Phonology, pp. 49–56.
Dupoux, E., Pallier, C., Sebastian-Gallés, N., Mehler, J., 1997. A destressing 'deafness' in French? J. Mem. Lang. 36, 406–421.
Dupoux, E., Pallier, C., Kakehi, K., Mehler, J., 2001. New evidence for pre-lexical phonological processing in word-recognition. Lang. Cognit. Process. 16 (5/6), 491–505.
Dupoux, E., Peperkamp, S., Sebastián-Gallés, N., 2001. A robust method to study stress ''deafness''. J. Acoust. Soc. Am. 110 (3), 1606–1618.
Dutoit, T., Pagel, V., Pierret, N., Bataille, F., van der Vreken, O., 1996. The MBROLA project: towards a set of high-quality speech synthesizers free of use for non-commercial purposes. In: Proceedings of The Fourth International Conference on Spoken Language Processing (ICSLP'96). 3.
Frisch, S.A., 1996. Frequency and Similarity in Phonology. PhD dissertation. Northwestern University.
Frisch, S.A., 2004. Language processing and segmental OCP effects. In: Hayes, B., Kirchner, R.M., Steriade, D. (Eds.), Phonetically Based Phonology. Cambridge University Press, Cambridge, pp. 346–371.
Frisch, S.A., Zawaydeh, B., 2001. The psychological reality of OCP-Place in Arabic. Lang. Acquis. 77, 91–106.
Frisch, S.A., Large, N.R., Pisoni, D.B., 2000. Perception of wordlikeness: effects of segment probability and length on the processing of nonwords. J. Mem. Lang. 42, 481–496.
Frisch, S.A., Large, N.R., Zawaydeh, B., Pisoni, D.B., 2001. Emergent phonotactic generalizations in English and Arabic. In: Bybee, J., Hopper, P. (Eds.), Frequency and the Emergence of Linguistic Structure. John Benjamins, Amsterdam, pp. 159–180.
Frisch, S.A., Pierrehumbert, J.B., Broe, M.B., 2004. Similarity avoidance and the OCP. Nat. Lang. Linguist. Theory 22 (1), 179–228.
Gallagher, G., 2013. Learning the identity effect as an artificial language: bias and generalisation. Phonology 30 (2), 253–295.
Gallagher, G., 2014. Evidence for an identity bias in phonotactics. Lab. Phonol. 5 (3), 337–378.
Gathercole, S.E., Frankish, C.R., Pickering, S.J., Peaker, S., 1999. Phonotactic influences on short-term memory. J. Exp. Psychol.: Learn. Mem. Cognit. 25, 84–95.
Goddijn, S., Binnenpoorte, D., 2003. Assessing manually corrected broad phonetic transcriptions in the Spoken Dutch Corpus. Paper Presented at the 15th ICPhS.
Graff, P., 2012. Communicative Efficiency in the Lexicon. PhD thesis. MIT.
Graff, P., Jaeger, T.F., 2009. Locality and feature specificity in OCP effects: evidence from Aymara, Dutch, and Javanese. In: Proceedings from Main Session of the 45th Meeting of the Chicago Linguistic Society.
Greenberg, J., 1950. The patterning of root morphemes in Semitic. Word 5, 162–181.
Hay, J., Pierrehumbert, J.B., Beckman, M., 2004. Speech perception, well-formedness, and the statistics of the lexicon. In: Local, J., Ogden, R., Temple, R. (Eds.), Phonetic Interpretation: Papers in Laboratory Phonology VI. Cambridge University Press, Cambridge UK, pp. 58–74.
Hayes, B., 1999. Phonetically-driven phonology: the role of optimality theory and inductive grounding. In: Darnell, M., Moravscik, E., Noonan, M., Newmeyer, F., Wheatly, K. (Eds.), Functionalism and Formalism in Linguistics, 1. John Benjamins, Amsterdam, pp. 243–285.
Hayes, B., 2000. Gradient well-formedness in optimality theory. In: Dekkers, J., Leeuw, F.v.d., Weijer, J.v.d. (Eds.), Optimality Theory: Phonology, Syntax, and Acquisition. Oxford University Press, Oxford.
Hayes, B., Wilson, C., 2008. A maximum entropy model of phonotactics and phonotactic learning. Linguist. Inq. 39 (3), 379–440.
Jaeger, T.F., 2008. Categorical data analysis: away from ANOVAs (transformation or not) and towards logit mixed models. J. Mem. Lang. 59, 434–446.
Kager, R., Shatzman, K., 2007. Phonological constraints in speech processing. In: Los, B., Koppen, M.v. (Eds.), Linguistics in the Netherlands 2007. John Benjamins, Amsterdam, pp. 100–111.
Leben, W., 1980. Suprasegmental Phonology. Garland, New York.
Luce, P.A., 1986. Neighborhoods of words in the mental lexicon. Doctoral dissertation, Indiana University, Bloomington, IN.
MacEachern, M., 1999. Laryngeal Co-Occurrence Restrictions. Garland, New York.
Mayer, T., Rohrdantz, C., Plank, F., Bak, P., Butt, M., Keim, D.A., 2010. Consonant co-occurrence in stems across languages: automatic analysis and visualization of a phonotactic constraint. Paper Presented at the Proceedings of the ACL 2010 Workshop on NLP and Linguistics: Finding the Common Ground.
McCarthy, J.J., 1985. Formal Problems in Semitic Phonology and Morphology. Garland, New York.
McCarthy, J.J., 1986. OCP effects: gemination and antigemination. Linguist. Inq. 17, 207–263.
McCarthy, J.J., 1988. Feature geometry and dependency: a review. Phonetica 45, 84–108.
McQueen, J., 1998. Segmentation of continuous speech using phonotactics. J. Mem. Lang. 39 (1), 21–46.
Norris, D., McQueen, J., Cutler, A., 1995. Competition and segmentation in spoken-word recognition. J. Exp. Psychol.: Learn. Mem. Cognit. 21, 1209–1228.
Padgett, J., 1995. Stricture in Feature Geometry. CSLI Publications, Stanford.
Peña, M., Bonatti, L., Nespor, M., Mehler, J., 2002. Signal-driven computations in speech processing. Science 298 (5593), 604–607.
Pierrehumbert, J.B., 1993. Dissimilarity in the Arabic verbal roots. Paper Presented at the Proceedings of the North East Linguistics Society.
Pierrehumbert, J.B., 2003a. Probabilistic phonology: discrimination and robustness. In: Bod, R., Hay, J., Jannedy, S. (Eds.), Probabilistic Linguistics. MIT Press, Cambridge, MA, pp. 177–228.
Pierrehumbert, J.B., 2003b. Phonetic diversity, statistical learning, and acquisition of phonology. Lang. Speech 46 (2-3), 115–154.
Plaut, D., McClelland, J., Seidenberg, M., Patterson, K., 1996. Understanding normal and impaired word reading: computational principles in quasi-regular domains. Psychol. Rev. 103, 56–115.
Pozdniakov, K., Segerer, G., 2007. Similar place avoidance: a statistical universal. Linguist. Typol. 11, 307–348.
R Core Team, 2012. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.
Scholes, R., 1966. Phonotactic Grammaticality. Mouton, The Hague.
Storkel, H.L., 2001. Learning new words: phonotactic probability in language development. J. Speech Lang. Hear. Res. 44, 1321–1337.
Suomi, K., McQueen, J., Cutler, A., 1997. Vowel harmony and speech segmentation in Finnish. J. Mem. Lang. 36 (3), 422–444.

Vitevitch, M.S., Luce, P.A., 1998. When words compete: levels of processing in perception of spoken words. Psychol. Sci. 9 (4), 325–329.
Vitevitch, M.S., Luce, P.A., 1999. Probabilistic phonotactics and neighbourhood activation in spoken word recognition. J. Mem. Lang. 40, 374–408.
Vroomen, J., Tuomainen, J., de Gelder, B., 1998. The roles of word stress and vowel harmony in speech segmentation. J. Mem. Lang. 38 (2), 133–149.
Wickham, H., 2009. ggplot2: Elegant Graphics for Data Analysis. Springer, New York.
Wilson, C., Obdeyn, M., 2009. Simplifying subsidiary theory: statistical evidence from Arabic, Muna, Shona, and Wargamay. Unpublished MS. Johns Hopkins University.
Yip, M., 1989. Feature geometry and cooccurrence restrictions. Phonology 6 (2), 349–374.