The Twenty-Fourth Manchester Phonology Meeting

ABSTRACTS BOOKLET

Thursday 26th - Saturday 28th May 2016

Held at
Hulme Hall, Manchester

Organised by a collaboration of phonologists at the **University of Edinburgh**, the **University of Manchester**, and elsewhere.

This booklet contains the abstracts for all the papers presented at the **twenty-fourth Manchester Phonology Meeting**, held at Hulme Hall, Manchester, in May 2016.

The abstracts are arranged in alphabetical order by the surname of the (first named) presenter. If any abstracts are missing from this booklet, it is most likely because the authors did not submit a non-anonymous version of their abstract.

The abstracts for the **oral paper sessions** are presented first, followed by the abstracts for the **poster paper sessions**, and the booklet concludes with abstracts for the **special session**.

The **final programme**, included in your registration pack and available on the conference website, gives the details of which papers are in which room, and at which times.

Oral papers

Modelling the Production and Perception of Kaqchikel Mayan with 'Small' Data

Ryan Bennett, Yale University, ryan.bennett@yale.edu; Kevin Tang, Yale University, kevin.tang@yale.edu

Motivation: No linguistic research is completed without understanding the interplay between production, perception and lexicon. The age of 'big data' has facilitated large-scale research on the effects of lexical statistics on phonological and phonetic patterning. Specially, recent work has shown that lexical factors such as functional load, distributional overlap and context-free/sensitive frequency, play a non-trival role in production and perception (Wedel et al. 2013, Seyfarth, 2014; Currie-Hall & Hume, 2015). However, it remains true that the majority of the world's languages are under-resourced, being spoken by relatively small populations. This renders large-scale lexical research extremely difficult, if not impossible, for most lesser-studied languages (Anand et al. 2011). Using the under-resourced language Kaqchikel (Mayan) as a test case, we provide a proof-of-concept that psycholinguistic and phonetic norms computed from less-than-ideal corpora are sufficient for predicting fine-grained differences in production and perception, such as the confusability of stops and word duration.

Methods: To model the lexical influence on production and perception of Kaqchikel, we compiled two new corpora, one written and one spoken, in order to extract lexical and acoustic variables; furthermore a perception study on stops was conducted. The **written** corpus was constructed from existing religious texts, spoken transcripts, government documents and educational books. The resulting corpus contains 1 million word tokens. This is far from ideal, in terms of its reliability for estimating low-frequency words (Brysbaert and New 2009), and its register (not speech-like). The **spoken** corpus consists of roughly 4 hours of spontaneous spoken Kaqchikel, collected in Guatemala in 2013 by one of the authors. The recordings are free narratives, told as monologues by 16 native speakers of Kaqchikel (19-84 years old, mean=33, median=28; 6 male). The corpus is not large: it includes only 40,000 or so word tokens. A subset of the corpus (~1 hour) was submitted to forced alignment (Gorman et al. 2011) for the purposes of acoustic analysis. The **perception** study was designed to focus on stops (plain and ejective stops). Stops are embedded in the CV and VC templates with the vowels [a, i, u], and they were spoken by a native Kaqchikel speaker. Speech-shaped noise was added to the stimuli at SNR of 0dB. The stimuli were presented in pairs in an AX-task to ~45 native hearers, and the accuracy and reaction time of each trial were recorded. To the best of our knowledge, these three datasets are the biggest of their kind for Kaqchikel.

Analyses: Two analyses were conducted over these three datasets:

Perceptual similarity: To model the acoustic influence on perceptual similarity, we extracted acoustic parameters from the spoken corpus such as phonation and spectral information of the stops and their adjacent vowels. These were then subjected to classification models (e.g. Random Forest) that are suitable for comparing correlated variables. This allowed us to identify the relative importance of each acoustic cue in perception. Crucially, we found that in addition to the acoustic influence, different lexical factors seem to play a role such as functional load, context-free and context-sensitive entropy measures.

Word frequency effect: We conducted a corpus-based analysis of durational reduction in spontaneous Kaqchikel speech. It has long been observed that words are often reduced in production when they are contextually predictable. This reduction takes a range of forms, including word/syllable shortening, segment/syllable deletion, and others (Hooper 1976; Whalen 1991; Aylett and Turk 2006). Focusing on the overall predictability of a word, more frequent words are more predictable than less frequent words, therefore they should also be shorter in duration. Token frequency was extracted from the text corpus. After controlling for word length, phrase-final lengthening, speaker- and sentence-level variation, Token Frequency remained a highly significant predictor of word duration ($B = -0.104$, $t = -7.01$, $p < 0.001$).

Conclusion: In summary, we established a proof-of-concept for under-resourced languages: it is possible to model the effect of lexical factors on phonetic and perceptual variation with less-than-ideal spoken and text corpora. Specifically, our analyses of Kaqchikel highlighted that (a) perceptual similarity is a function of acoustic cues as well as lexical statistics; crucially these factors are weighted unequally – that is not all lexical statistics and acoustic cues are used by the perceptual system; (b) the classic word frequency effect on duration remains robust even in a sparse corpus, and in a highly-inflecting, morphologically-complex language like Kaqchikel.

The typology of voiced spirants and the emergence of [sonorant]

Christina Bjorndahl, cbjorn@andrew.cmu.edu
Carnegie Mellon University and Cornell University

It is well known that for obstruents, voicelessness is the default, a fact often phrased as an implicational relation: if a language has a voiced obstruent, it has its voiceless counterpart (Maddieson, 1984). For stops and sibilants, the implicational relation is robust: among languages in the UPSID database, only 1.1% have voiced stops without corresponding voiceless stops, and none have a voiced sibilant without a voiceless sibilant. For the non-sibilant fricative pairs {[ɸ, β], [f, v], [θ, ð], [x, ɣ]} however, *rates of violation* are far higher, ranging from 22% for labiodental [f, v] to 75% for bilabial [ɸ, β]. Such high rates of violations within the non-sibilant fricatives (henceforth *spirants*) suggest that the relationship of voiced to voiceless spirants does not parallel that of other voiced-voiceless obstruent pairs. We argue that the typological facts, phonological patterning, and phonetic character of voiced spirants support an *emergent* interpretation of the [sonorant] feature (Mielke, 2008).

Drawing from a database study of the UPSID languages, our first objective is to show that the typological differences between voiced sibilants and spirants are more profound than previously appreciated. For example, Maddieson (2010) observes that if a language has a voiced fricative, it is likely to have a voiced stop; we show that this implicational relation is driven entirely by the voiced sibilants, and in fact fails for the voiced spirants. Therefore, while voicing in sibilants is, in some sense, parasitic on stop voicing, spirant voicing is not. Moreover, voiced sibilants and spirants show a different distribution according to inventory size: Lindblom and Maddieson (1988) argue that the likelihood of an inventory containing *elaborated articulations*, e.g., voiced fricatives, increases proportionally with inventory size; again, this is true for voiced sibilants, but not for voiced spirants. Furthermore, we show that [β, v, ð, ɣ] are most likely to occur unpaired (i.e., without [ɸ, f, θ, x], respectively) in small inventories, but rarely occur unpaired in large inventories.

Due to the phonetic character of voiced spirants, namely, the aerodynamic tensions that arise from maintaining voicing and frication, voiced spirants make for poor obstruents (Ohala, 1983). This fact, combined with their sonorant patterning in various languages, leads Botma and van't Veer (2013) to conclude that in many cases such segments have been mistranscribed, and are better classified as sonorants. Such an analysis is attractive because it leaves the implicational relations with respect to obstruent voicing intact. If unpaired [β, v, ð, ɣ] are in fact sonorants, then it is not surprising that they surface without voiceless counterparts and independently of whether there are voiced stops in the inventory. However, while a [+sonorant] classification may be appropriate for any given language, we argue that it does not go far enough in elucidating the subtlety in the cross-linguistic identity of the voiced spirants. From a typological perspective, such an analysis does not explain the interaction between implicational relations and inventory size. Phonologically, just as voiced spirants are poor obstruents, they are also ill-suited to sonorant membership: they are rarely syllabic, have restricted phonotactics that do not parallel other sonorants, and often exhibit at least partial phonological pairing with their voiceless counterparts. Furthermore, voiced spirants can pattern with sonorants despite the presence of their voiceless counterpart, such as Serbian /v/ (Morén, 2003), which is not addressed by Botma and van't Veer (2013). Finally, voiced spirants rarely contrast with approximants at the same place of articulation (e.g., /v/ ∼ /ʋ/), but when they do, such contrasts appear to be accompanied by additional parameters of phonetic enhancement (Stevens and Keyser, 1989), as in Dutch (Gussenhoven and Bremmer Jr., 1983).

Building on Dresher (2014), we propose that [sonorant] is a *phonological* construct that is induced in the course of acquisition. This, together with the model of *phonetic* inventory dispersion proposed by Lindblom and Maddieson (1988), provides a framework well suited for explaining the typological intermediacy, phonological ambivalence, and phonetic realization of the voiced spirants with respect to the obstruent-sonorant divide.

Opaque domain modelling in OT: against stepwise prosodic parsing in Harmonic Serialism

Karolina Broś, University of Warsaw (k.bros@uw.edu.pl)

One of the recent concerns in OT research is the way prosodic structure is built and to what extent it interacts with phonological processes. In the framework of Harmonic Serialism (McCarthy 2008), various attempts have been made at incorporating prosody into stepwise evaluation in order to accommodate mappings generated in languages that permit e.g. resyllabification and ambisyllabicity (e.g. Elfner 2009, Pater 2010). Torres-Tamarit's (2012) model seems to especially efficient in dealing with various types of overapplication, such as Spanish opaque aspiration across different prosodic domains. According to this model, prosodic parsing takes place in steps across morphological constituents. Thus, syllabification applies to each morph separately, while morphophonological alignment ensures coincidence of morphological and prosodic edges to the extent of a given language's syllabification properties and misalignment permitted across phonological words (e.g. in prefixation and phrase level phonology). More specifically, to prevent the projection of a core syllable across a prosodic domain, two alignment constraints are needed: ALIGN-L(stem, PW) and ALIGN-L(MWd, PW). At the same time, the morphosyntax-prosody mapping is effected via parsing (PARSE-SEG, PARSE-PW). The interaction of the two types of constraints with markedness constraints banning illicit structures makes it possible to account for opacity effects observed in domains created by prosodic constituents. Such a model is capable of generating a wide range of dialectal differences involving across-prefix and phrase-level opaque aspiration (e.g. opacity across all domains: *desarmar* 'to disarm' [de.haɾ.maɾ], *las heridas* 'the wounds' [la.he.ɾi.ðah], opacity only across a word boundary: [de.saɾ.maɾ], [la.he.ɾi.ðah], or no opacity: [de.saɾ.maɾ], [la.se.ɾi.ðah]). Nevertheless, certain mappings fail to be generated by this model. When the opaque interaction involves two competing repairs, prosody ceases to be the deciding factor. In this paper, I argue that stepwise prosodic parsing is insufficient for explaining certain opacity effects. I demonstrate that cyclic domain construction by means of strata is more suitable for explaining prosody-morphology interactions. Stratal OT (Bermúdez-Otero, to appear) is superior to HS in that it provides a straightforward explanation of opaque interactions by assigning phonological processes to different domains. The Chilean dialect of Spanish (Broś 2012) is a case in point. While it presents both deletion and debuccalisation of coda *s* in word-final position (prepausal/preconsonantal, and prevocalic, respectively), it is far less radical in the treatment of prefix-final *s*. Thus, the observed outputs are: *desarmar* [de.**s**aɾ.maɾ], and *las heridas* [la.**h**e.ɾi.ða**Ø**]. In line with Torres-Tamarit's model, in any sequence of words, each lexical item is a separate morphological word and has to be parsed into a prosodic word separately. Only then is it prosodified further into a phonological phrase, and resyllabification across prosodic constituents is enabled. In an HS evaluation, the ranking ALIGN-L (stem, Pwd)>>PARSE-SEG>>ONSET>>*S]CODA>> PARSE(PWD)>>IDENT(Pl)>>MAX(Seg), ALIGN-L (STEM, PWD) ensures that prosodic words are parsed before syllabification at Step 1 ([las]#[eɾiðas]. This is followed by syllabification at Step 2: [(las)]#[(e)(ɾi)(ðas)] by the high-ranked PARSE-SEG. At Step 3, *S]CODA mandates a repair before the two words are parsed into a phonological phrase due to ALIGN-L (MWD, PWD), but the ranking will choose deletion as a better option than debuccalisation ([(la)]#[(e)(ɾi)(ða)]). Once the first coda *s* is deleted, there is no way of restoring it at a later stage without violating higher-ranked constraints. Thus, sophisticated prosodic parsing does not prevent unattested mappings from being generated in opaque domains involving two competing repairs. Against this background, the cyclicity model offered by Stratal OT is superior to HS in one crucial respect: its very design assumes serial prosodic structure building associated with morphophonological strata, and each stratum is based on different phonological predictions, hence a different ranking. By assuming that the discussed processes apply at different strata, the opacity problem is solved. What is more, constraints related to prosodic structure also differentiate between word-internal position and word edges. By invoking contiguity inside prosodic words (CONTIGPW), a Stratal OT analysis is able to account both for the advancement of *s* weakening at word edges and for the retention of coda *s* inside words. Morpheme internal *s* is aspirated but protected from deletion, and prefix-final *s* is retained and resyllabified directly with the following stem (and hence protected in the domain of PW').

Characterising the black sheep of phonology: a unified account of French & English sC clusters

Faith Chiu (*University College London*; faith.chiu.11@ucl.ac.uk) & Typhanie Prince (*University of Nantes*)

Introduction. It is well known that word-initial s+obstruent clusters (sT) violate the Sonority Sequencing Principle (Selkirk, 1984; Clements, 1990). The SSP stipulates that left-edge well-formedness is defined by a rise in sonority from onset(s) to the peak nucleus. Exceptionally, in sT, this cluster-initial /s/ is more or equally as sonorous as its following obstruent. Thus, such initial sT clusters (like those in French and English) have been analysed in a plethora of ways as not being a "true" branching onset. Besides the contour segment interpretation (Steriade, 1982), most analyses classify cluster-initial /s/ in sT as special. Along the general vein of extraprosodic (Goldsmith, 1990) or extrasyllabic claims (Halle & Vergnaud, 1980), cluster-initial /s/ has also been labelled as an appendix (Vaux & Wolfe, 2009), or, where /s/ employs an existing constituent, it is a coda in a coda-onset sequence (Kaye, 1992; Goad, 2012), or syllabic (Seigneur-Froli, 2006; Barillot & Rizzolo, 2012).

Research question; outline. First, we evaluate the above constituency approaches against original experimental data. Predictions made adopting the syllabic s analysis (Seigneur-Froli, 2006) best explain phonological behaviour in our French production, and English perception and processing results. In this syllabic s approach, /s/ branches onto the following empty nucleus. We then propose a unified account for initial, medial and final French and English sC clusters (s+obstruent and s+sonorant) using a CVCV framework (Lowenstamm, 1996; Scheer, 2004).

Methodology. Clinical and L1 acquisition data for French production is obtained through naming and repetition tasks in aphasic patients (n=20) and normally developing children (n=20, mean age=2;3). English perception and processing results come from a dichotic listening experiment and an EEG task analysed for event-related potentials (ERP) in British English adults (n=10).

Results. Initial sC are identical to obstruent-obstruent clusters (TT) in behaviour (but not in its representation due to its existence in #TR-only languages, as we show). Initial English s+sonorant clusters resemble sT and not branching onsets (*contra* Giegerich, 1992). Medial sT are identical to TT in behaviour (and in representation). Final sT differ from TT in behaviour (and in representation).

Analysis. Following Scheer (2004), Brun-Trigaud & Scheer (2010) and Scheer & Zikova (2010), we propose these representations for the initial **#**sC (1a,b), medial **V**sC**V** (1c) and final sC**#** (1d); initial = final. French and English possess an initial CV (Lowenstamm, 1999; Scheer, 2004).

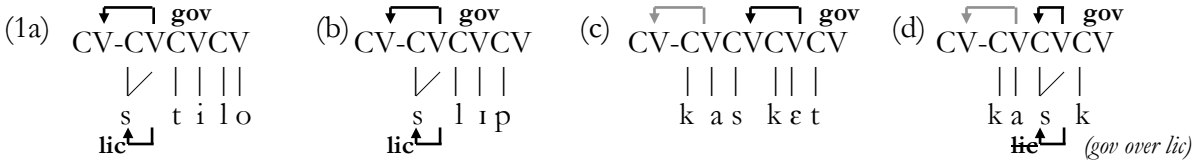

When considering the empty nucleus (EN) *within* the clusters, the difference between #sC and TT is government: the EN in #sC is ungoverned. /s/ branches onto the following V to provide lateral actorship for this nucleus. When initial, s-branching is crucial to explain its well-formed status in a #TR-only language. Otherwise, either the initial CV or the enclosed EN in #sC would be an orphan. This theoretically derived distinction in representation is identically demonstrated in behaviour. In our English data, #sC, like TT, shows no perceptual fusion in dichotic listening: no branching onsets (TR) formed; infrasegmental government (IG) fails. Yet in EEG, #sC has no ill-formedness ERP: it is unlike English *#TT that has 2 successive ungoverned EN. Extending the coda mirror (Ségéral & Scheer, 2001), final sC# also branches, making it *only* governed with government-over-licensing as in S&Z's 'intervocalic' site. In our French aphasic data, repairs for sC# differ from ʁC#. C1 in **ʁ**C# but not **s**C# can be deleted, but C1 can be deleted in both V**s**CV and V**ʁ**CV. Thus, only ungoverned and unlicensed segments are deleted. Our analysis also shows C1 in both #**s**C and #**T**R are ungoverned but licensed (*c.f.* Kaye, 1992). We explain #sC/#TR differences at melody (hinted by Scheer, 2004): lack of dichotic fusion and production deletions of /s/ in #sC is due to the lack of IG, not licensing.

Sources of phonological variation:
Evidence from a naturalistic corpus of Jakarta Indonesian

Abigail C. Cohn, Cornell U. acc4@cornell.edu
Ferdinan Okki Kurniawan, Cornell U. fok3@cornell.edu

Increased incorporation of experimental and naturalistic data in phonological analysis in recent years has provided an important complement to introspection and elicited data. These data sources highlight that there is greater gradience (temporal/spatial) and variability (variation in frequency of occurrence) in phonological and morphophonological patterning than often thought. While it is indeed a positive development that increasingly attention to such variation is being incorporated into phonological analyses, some caution is also in order in terms of understanding the relevant contributions of phonological, lexical, and inter- and intra speaker factors. To be useful for phonetic and phonological analysis, phonologists need naturalistic corpora that are phonetically transcribed, linked to corresponding audio files, and enriched with sufficient metadata to tease apart inter- and intra-speaker variation.

In this talk we discuss two examples from Indonesian, highlighting the interplay between phonological patterning and both inter- and intra- speaker variation. Most prior work on Indonesian has focused on Standard Indonesian. However, as a formal standard variety, Standard Indonesian (SI) is subject to normative aspects of pronunciation that affect observed patterns. To address these limitations, we investigate observed phonological and morphophonological patterning in Jakarta Indonesian (JI), a rapidly developing colloquial variety spoken in and around the capital. We draw on data from a phonetically transcribed naturalistic spoken corpus of Jakarta Indonesian (Gil and Tadmor 2015) along with acoustic analysis of corresponding audio files.

First, we investigate the realization of the active prefix *N-* in JI (cognate with the *meN-* prefix in SI). It surfaces as [ŋ(ə)-] with root-initial approximants and vowels (/N+lamar/ [ŋəlamar] 'to propose'), and "nasal substitution" occurs with initial voiceless stops and nasals (/N+pilih/ [milih] 'to choose'); however with voiced stop-initial roots, two variants are observed: /N+bəli/ [ŋəbəli] ~ [mbəli] 'to buy'. Evidence from corpus data and a production study both attest to extensive inter- and intra-speaker variation, suggesting at least a partially historical explanation. Some speakers primarily use one or the other pattern, while some speakers indeed use both forms, and not predictable from specific item or place of articulation of the root initial consonant. Thus we do not observe lexicalization of specific forms (such as was observed by Zuraw 2010 for a similar pattern in Tagalog).

Second, we investigate the status of schwa based on data from the same corpus. While schwa in Indonesian is contrastive (*gamelan* [gaməlan] 'musical ensemble' ~ *jumlah* [jumla] 'total'), the distribution of schwa is different from the other five vowels in terms of stress and phonotactics, and the realization of schwa is described to be variable. We conclude that synchronically the variability arises from optional deletion of underlying schwa, not insertion of schwa to break up historically unallowable clusters. Furthermore schwa is not gradient in its realization; it is either present or absent (suggesting it is not a by-product of gestural alignment). Thus it is variable, but not gradient in its realization. We show that multiple factors together condition the optional deletion of schwa, including orthography, historical source, morphological structure, and phonological structure.

These data from a naturalistic corpus highlight the extent and complexity of variation seen in phonological and morphophonological patterns. Much can be learned from such corpus work, but careful analysis is needed. Multiple factors including phonological environment, morphological structure, speaker identity, historical source, encoded in part in the orthography, may all play a role. Each of these factors needs to be encoded in the corpus and included in the analysis to fully understand these complex interactions.

Evidence for unmarked consonants from loanword adaptation

Laura J. Downing, Göteborgs universitet

Vowel epenthesis is a well-documented strategy for adapting loanwords into languages that have less complex syllable structure than the donor language(s). For example, Uffmann (2006, 2007) shows that when English words are borrowed into Shona, consonant clusters (except NC) and word-final consonants – both violations of Shona's phonotactics – trigger vowel epenthesis: ***k**i**r**a**k**i* 'clerk'. In examples like this, the epenthetic vowel appears to be a fixed, default vowel, 'i'. Indeed, loanword epenthesis is a commonly cited source of evidence for choosing the default vowel of a language (N. Hall 2011).

However, work like Rose & Demuth (2006), Uffmann (2006, 2007) shows that the choice of epenthetic vowel is often more complex than this. Two other strategies are widely attested: the epenthetic vowel can agree in features with a preceding consonant, e.g., Shona *kira<u>b</u>h**u*** 'club'; or the epenthetic vowel can copy a vowel in an adjacent syllable, skipping an intervening consonant, e.g., Shona *tiro<u>k</u>**o*** 'truck' (Uffmann 2006). The same language can use all three strategies, with the choice depending on the context.

<u>The issue</u>: Rose & Demuth's and Uffmann's studies show that, even among closely related languages, like Bantu languages, one finds variation in the choice of strategy in particular phonological contexts. As they argue, this provides strong evidence that vowel epenthesis in loanwords is phonological, not just phonetic. The challenge is how to account for the range of variation.

<u>The problem</u>: Uffmann (2006, 2007) notes two asymmetries involving velar obstruents in the large, multilingual corpus of loanwords that he investigates. The epenthetic vowel never agrees in place with a preceding velar (dorsal) consonant: e.g. ***k**i**r**a**bh**u*** 'club'. Only coronals and labials trigger agreement. Also, velar stops are often transparent for vowel copy, as in *tiro**k**o* 'truck', whereas consonants with other places of articulation trigger vowel agreement. Uffmann accounts for these asymmetries by appealing to two harmonic scales:

(1) *LINK(C,V)/DOR >> *LINK(C,V)/LAB >> *LINK(C,V)/COR

(2) *SKIP(V) >> *SKIP(SON) >> *SKIP(FRIC) >> *SKIP(STOP) >> *SKIP(LAR)

Together they penalize linking dorsal features with a CV sequence and favor skipping a stop in vowel copy.

Since the constraints in (1) and (2) are harmonically ranked scales, they predict that one should not find a language where both – and only – a velar obstruent stop and a coronal sonorant are productively transparent for vowel copy in loanword adaptation. These two consonant types are at the opposite end of the harmonic scales. Yet, that is exactly what we find in Chichewa, a Bantu language spoken in Malawi.

As in Shona, we find three vowel epenthesis strategies. Labial and coronal consonants (except the liquid 'l') trigger agreement on a following epenthetic vowel: *b**u**leki* 'brake', *kap**u**ten**i*** 'captain'; *haf**u*** 'half'; *s**i**tov**u*** 'stove'. Epenthetic vowels following dorsals are either default 'i' or show vowel copy, as expected in Uffmann's analysis: *buk**u*** 'book'; *k**a**lasi* 'class'; *bok**o**si* 'box'; *Ag**i**nesi* 'Agnes'; *wig**i*** 'wig'. However, the liquid 'l' – uncontroversially a coronal sonorant – is almost always (preceded or) followed by an epenthetic vowel that copies the vowel of an adjacent syllable: *ci-bangil**i*** 'bangles, bracelet'; *d**a**layival**a*** 'driver'; *sitol**o*** 'store'; *kol**o**geti* 'Colgate'; *sukul**u*** 'school'; *d**e**lesi* 'dress'.

What I will argue is that velar stops and the liquid 'l' do not trigger place agreement on an epenthetic vowel and allow vowel copy because these are the least marked consonants of Chichewa. Only marked consonants spread place features, and only unmarked consonants are skipped for vowel copy. In sum, loanword adaptation provides a window into the least marked consonants, as well as to the least marked vowels, of a language.

Latinate Suffixes and the Directionality of English Stress

B. Elan Dresher, University of Toronto & Aditi Lahiri, Oxford University
dresher@chass.utoronto.ca aditi.lahiri@ling-phil.ox.ac.uk

A major change in the history of English stress was a shift in directionality (edge orientation) from the Germanic practice of computing stress from the left edge of the word (apart from certain unstressed prefixes) to the right edge (Halle & Keyser 1971). All writers on the topic attribute this shift to the massive influx of Romance loanwords into English, beginning with the Norman Conquest (1066), but there has been no agreement as to when the shift began. Some writers (Halle & Keyser 1971; Lass 1992) have proposed that the Romance right-edge oriented stress rule gained a foothold in English in the Middle English period of Chaucer (c1343–1400), though it did not become the main stress rule until some time later. We take the view of those who argue that this wave of loanwords, mostly via French, did not have any lasting impact on English stress (Jordan 1974: 199; Minkova 1997, 2006; Redford 2003).

Danielsson (1948) and Poldauf (1981) associate the change in directionality with the accumulation of words with Latin and French suffixes such as *-able/-ible, -ation, -ator, -ic(al), -ity,* etc. In such words, stress is computed from the right: for example, *medícinal* and *philosóph-ical* can be assigned stress by a unified rule computing from the right edge (both have stress on the antepenult), but not from the left. In the earlier periods, Latin words had been borrowed as morphologically simplex (Minkova & Stockwell 1996; Lahiri & Fikkert 1999); it took a while for these borrowings to become common enough that English speakers could recognize their morphological composition. At that point they could identify recurring morphemes, such as derivational suffixes; the rightward directionality of stress in words with these suffixes could then become apparent. Moreover, the native English vocabulary did not have words of similar complexity that could systematically contradict these patterns.

When did this shift occur? We will investigate this question by looking at the data from a quantitative point of view. As an approximation to when a word first entered the language, we will use the earliest date for which it is attested in the *Oxford English Dictionary* (*OED*). We will compare the distribution and composition of Romance loanwords at two points: 1400, the year that Chaucer died, and 1570, the year of the publication of Levins' *Manipulus Vocabulorum*, a reverse dictionary that indicates the location of stress in many words. Lahiri (2015: 239) proposes that Levins 1570 gives evidence of a stress system that had recently undergone a transition from a left- to a right-edge orientation (main stress, however, had not yet changed from left to right).

According to the *OED*, by 1400 English had borrowed 6,580 words of Romance (mostly French and Latin) origin, 21.5% of all words. By 1570, the Romance words had increased by 93%, but their percentage in the English vocabulary was 18.3%, *less* than it had been in 1400. However, the *composition* of the Romance vocabulary changed more drastically, showing a great increase in complex words with suffixes. From 1400 to 1570, the number of words with *-able* went from 204 to 906, an increase of 344%; words with adjective-forming *–al* increased from 163 to 745 (357%); and words with *-ity* increased from 144 to 563 (291%). These numbers lend *prima facie* support to the notion that the relevant data for learners in 1570 was very different from what it had been in 1400. We will consider this data in the light of Yang's (2005) Tolerance Principle, that exceptions to a rule can be tolerated if they are less than or equal to $N/\ln N$ (we call this the Yang Threshold, Y). If, as a first approximation, we set N = the total number of words in the language at each period, we find that the number of words with relevant Latinate suffixes is only 60%Y in 1400, but 107%Y in 1570. These results support the hypothesis that Latinate loanwords with suffixes caused the shift in the directionality of English stress when their number exceeded Y, at some point before 1570.

Phonological Knowledge Beyond the Word-domain: The Case of Illusory Consonants

Karthik Durvasula, Chenchen Xu, Mingzhe Zheng, Xiaomei Wang, Yen-Hwei Lin
Michigan State University
{durvasul, xuchenc1, zhengmi4, wangxi89, liny}@msu.edu

Native listeners perceive illusory sounds, typically when presented with sound sequences that do not respect the *word-internal* phonotactic constraints of their language (Dehaene-Lambertz, Dupoux, & Gout, 2000; Dupoux, Kakehi, Hirose, Pallier, & Mehler, 1999). Furthermore, some have argued that the illusory segment can only be the phonetically minimal vowel in the language (Dupoux, Parlato, Frota, Hirose, & Peperkamp, 2011), while others have argued that the illusory segment depends on the phonology of the language (Durvasula & Kahng, 2016). Here we focus on *consonant-vowel* sequences, which are illicit word-internally in Mandarin Chinese and American English, but are licit outside the *word-domain* in American English, to show that even such contexts trigger illusory segments, and that the illusory segment (in this case, a glide) depends on the phonology of the language.

Following Durvasula & Kahng (2016), we suggest that the task of the listener in speech perception is to identify the licit phonemic representation that best maps to the given acoustic input. In Mandarin, alveopalatal consonants [tɕ, tɕ^h, ɕ] can never appear directly before a non-high front vowel; there must be front palatal glide [j, **ɥ**] between the alveopalatal and non-high front vowel in such contexts – so [tɕ^{h}jɑʊ, but *tɕ^{h}ɑʊ]. Given the acoustic cues present in a sequence such as [ɑtɕ^{h}ɑʊ], and the phonological patterns of Mandarin Chinese, the best phonemic parse of the input is one that includes the illusory /j/. Therefore, when Mandarin listeners are presented with a phonotactically illicit sequence involving an alveopalatal consonant (e.g., [ɑtɕ^{h}ɑʊ]), we predict that the sequence is repaired by the perception of an illusory /j/ between the alveopalatal consonant and the vowel. As a consequence, Mandarin speakers should have difficulty distinguishing [ɑtʃjɑʊ] from [ɑtʃɑʊ]. In American English, palato-alveolars [tʃ, ʃ, dʒ, ʒ] cannot appear before [j] within words, but they can appear adjacent to [j] across words or compounds. Therefore, for English listeners, [tɕ^h] is expected to be perceived as /tʃ/, and both the acoustic inputs [ɑtʃjɑʊ] and [ɑtʃɑʊ] have valid phonemic parses; crucially, [ɑtʃjɑʊ] has a valid phonemic parse as /ɑtʃ # jɑʊ/, though the sequence is disallowed within words (# = word-boundary). Therefore, the phonological patterns in English will allow English listeners to distinguish [ɑtʃjɑʊ] from [ɑtʃɑʊ]. As a result, we predict that English listeners should be able to discriminate between [ɑtɕ^{h}jɑʊ] and [ɑtɕ^{h}ɑʊ] better than Mandarin listeners.

We ran an ABX experiment on 20 native Mandarin listeners, and 19 native American English listeners as controls. We presented participants with pairs of nonce words of the form **aC$_1$C$_2$V$_1$ma** [C$_1$ = t^h / tɕ^h; C$_2$ = j / Ø; V$_1$ = ɑʊ / Ø]. All tokens had initial stress, and were recorded by a trained phonetician. We take confusability between pairs of words with and without the glide to suggest the induction of an illusory segment. A mixed ANOVA of the percentages of correct responses revealed, crucially, an interaction of word-pair by language [$F(5,185) = 7.10$, $p<0.0001$, $\eta_{gen}^2 = .10$]. Planned two-sample Welch-tests of correct responses between the two language groups revealed that [atɕ^{h}ɑʊma-atɕ^{h}jɑʊma] was the only word-pair with a difference between the two language groups, with Mandarin listeners performing worse than English listeners [$\text{Mean}_{\text{Eng-Mand}} = 16\%$, $t(35.74) = 3.52$, $p=0.001$]. For other word-pairs, the mean difference was less than 3%.

Our results suggest that phonological knowledge, including knowledge beyond the word-domain, is employed to both identify the locus of a perceptual repair in phonotactically illicit noncewords and to infer the illusory segment itself.

There is no allomorphy in the Italian definite article

N. Faust (SFL/Paris 8), N. Lampitelli (LLL/Tours) & S. Ulfsbjorninn (Lyon 3/SFL)

Italian definite determiners exhibit phonologically-contextual allomorphy (1), but the selection of allomorphs does not readily lend itself to an analysis of phonological optimization. For this reason, multiple underlying forms are generally assumed.

(1) Definite articles	Singular	Plural	
Type A word (also branching onsets)	il sakko	i sakki	'bag'
Type B word (also initial geminates)	lo ska:fo	ʎi ska:fi	'hull'

An analysis with one underlying form /ilo/ has been proposed for the singular articles by Larsen (1998). The plural articles remained unexplained. Indeed, why do CV-initial bases take [i] and CCV-initial bases take [ʎi]? Why and how could a [ʎ] be formed? In this talk, we will show that the plural articles, too, can be derived from the same underlying representation. All one needs to do is replace the regular singular marker /-o/ in Larsen's representation with the regular masculine plural marker /-i/. The analysis benefits linguistic theory by showing that this important case of allomorphy is in fact just phonology at work.

With Larsen, we assume the determiner is composed of a CVCV template, where the vocalic content is underlyingly floating: /^{i}l^{o}/ (2-3). Nuclei with only floating melody are treated as empty when p-licensed (consistent with GP/Strict CV assumptions for vowel zero alternations, Kaye at al. 1990); but the melody associates to the position when this is not the case. Thus, in our case the docking of floating melody to V_1 and V_2 ultimately depends on the status of V_3. In type A words (2) V_3 is contentful, ensuring that V_2 is p-licensed and remains empty, and /^{o}/ remains afloat. As a result V_2 cannot p-license V_1, and /^{i}/ must dock onto this position, yielding [il]. Conversely, in type B words V_3 is p-licensed (3), and so V_2 cannot be p-licensed and must associate to /^{o}/. V_2 in this case acts as a licensor for V_1, and so /^{i}/ remains afloat. We derive [lo].

(2) [il sakko]

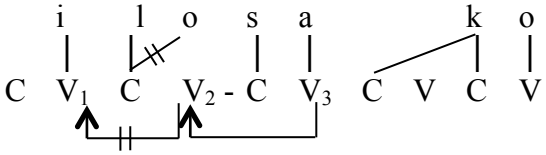

(3) [loska:fo]

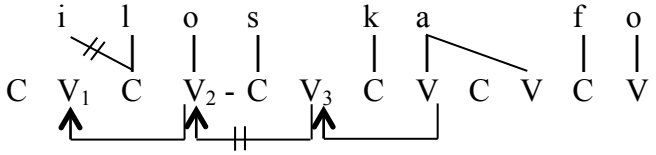

The plural definite article unsurprisingly has /i/ where the singular had /o/. Its UR is /^{i}l^{i}/. As we saw for the singular, either the first or the second floating /^{i}/ will remain afloat. In either case, this floating melody does not delete, but instead docks onto C_2, forming a unordered complex expression {i,l} or {l,i}. Universally, we expect this combination to yield [ʎ]. But in Italian this complexity is the only the *first* condition on forming the [ʎ]; there is an additional *language specific* condition, namely that {l,i} be <u>*licensed by a nucleus associated with the melody /i/*</u>. As shown in (4), this condition does not hold in type A words. Consequently the structure is invalid and the outcome is just [i]. In contrast, in type B words (5) V_2 is unlicensed and thus associated to the second floating /^{i}/. The complex expression {i,l} is licensed (←) by V_2 to derive [ʎi].

(4) [i sakki]

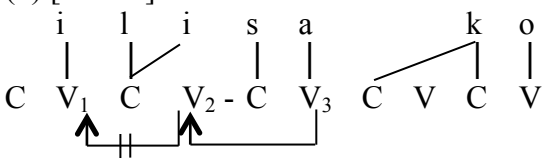

(5) [ʎi ska:fi]

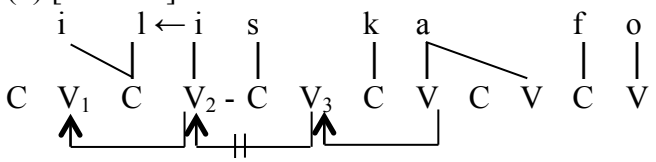

The analysis is then easily extended to the vowel-initial [l-azino] and [ʎj-a:zini] 'the donkey(s)', by assuming a constraint against hiatus, independently motivated in Italian. It is further shown that the proposal explains the 'inherent geminate' status of [ʎ] and other palatals in Italian (Chierchia 1986): in order to be palatalized, the segment always has to be preceded by a floating /i/ which comes with its own CV unit.

To summarize, looks can be deceiving: with the correct phonological analysis, even the Italian [i] vs. [ʎi] allomorphy is in fact regular phonology, not allomorphy: there is only one UR.

Simplicity versus data in phonological change

Josef Fruehwald
josef.frueh@ed.ac.uk
University of Edinburgh

Betsy Sneller
esnell@sas.upenn.edu
University of Pennsylvania

After a history of being at the nexus of a number of debates in phonology and sound change, (Labov, 1981; Kirparsky, 1993), the classic Philadelphia Short-a System (PHL) is being rapidly replaced by the nasal system (NAS) (Labov et al, 2015) which is considerably simpler to formally define. One reasonable hypothesis is that the shift from the complex PHL to the simple NAS was inevitable, since learners would naturally gravitate to the simpler grammar. However, on the basis of recent work on productivity and grammar competition by Yang (2000, 2002, 2005), **we will argue that the "simpler" NAS would not be a tenable grammar given PHL input data**. PHL is defined in (1) (tense before tautosyllabic anterior [nasals or voiceless fricatives]), and NAS in (2) (tense before nasals).

(1) æ → tense / __ {m, n, f, θ, s}]$_{\sigma}$]$_{stem}$ (2) æ → tense / __ {m, n, ŋ}

There are multiple tense and lax lexical exceptions to (1) and it is not a surface-true generalization, since it applies at the stem level. NAS, on the other hand, has tense /æ/ before all nasals regardless of syllabicity, making (2) a surface-true generalization.

We carried out both a productivity analysis, based on the Tolerance Principle from Yang (2005), and a grammar competition analysis based on Yang (2000, 2002), to evaluate the viability of (2) given input data from (1). The Tolerance Principle states that a rule can tolerate $T = N/\ln(N)$ lexical exceptions before it is more efficient to memorize all lexical items. The grammar competition model states that if G_2 can parse more data from G_1 than vice-versa, G_2 will replace G_1.

We coded N=2,258 word types from the CHILDES database (MacWhinney, 2000) for whether they would be tense or lax under PHL. We also coded the data according to versions of (1) which are surface true, lack the syllabic constraint, and lack any fricative conditioning, to test whether any version of PHL input data could be tolerated by a NAS grammar. **The only version of PHL data which passed $T = N/\ln(N) = 292$ for a NAS grammar was PHL−fricatives (266 exceptions).** We also trained two Noisy-HGs (Boersma & Pater, 2007): one on PHL data, one on NAS data. We then evaluated the accuracy of these Noisy-HGs when parsing data from the other grammar. The accuracy of PHL parsing NAS data was 83%; the accuracy of NAS parsing PHL data was 79.7%. The difference is small, but according to Yang (2002) **this means PHL ought to win in grammar competition with NAS**.

We conclude from these analyses that **NAS was not destined to replace PHL due to its formal simplicity**. The fact remains, however, that NAS *is* replacing PHL, which should be anomalous given these results. One possible explanation for this is that there is an independent change for speakers to *phonetically* lower their tense allophone (Labov et al, 2013), potentially leading to misperception of /æ/ in the pre-fricative contexts, which tend to be the least phonetically tense (Labov, 2001), as lax. This misperception would effectively be the PHL−fricatives grammar as input data, which was the only version of PHL which passed *T* for the NAS grammar. If this analysis is correct, it provides an example of a phonetic change altering the quality of phonological evidence and affecting phonological change.

Computing segmental and suprasegmental information in lexical decision

Guilherme Duarte Garcia (McGill University)
guilherme.garcia@mail.mcgill.ca

Studies on lexical decision have shown that factors such as word frequency and neighbourhood density significantly impact speakers' reaction time (RT) when deciding whether a word is real or not (e.g., Gardner et al. 1987, Cutler & Butterfield 1992, Vitevitch & Luce 1998). Phonological wellformedness has also been shown to affect lexical decision: more frequent phonotactic patterns often correlate with faster RTs (Bailey & Hahn 2001). Few studies, though, have looked at the effect of suprasegmental factors such as stress, the focus of the present study. For example, Vitevitch et al. (1997) show that word-initial stress tends to result in faster RTs for English speakers. However, given the positional bias towards word-initial stress in the English lexicon (Cutler and Carter 1987, Cutler and Norris 1987) and the fact that information near the left edge of words makes it a better cue for word recognition (Horowitz et al. 1968, 1969), it is not possible to accurately determine the reason why earlier stress correlates with faster RTs in English. For that, one needs a language with no bias towards word-initial stress.

The present paper investigates how stress affects lexical decision time in Portuguese, a language where three stress positions are available: antepenult (APU), penult (PU) or final (U). In Portuguese, unlike English, APU stress is the least frequent pattern, followed by U and PU stress (Bisol 1994, Wetzels 2007). If more frequent stress patterns correlate with faster lexical decision (mirroring word frequency effects), words with APU stress should yield longer RTs than words with PU and U stress in Portuguese. In this study, I show that the opposite is true, and propose an alternative explanation for the effects of stress on lexical decision. The explanation stems from the point at which one computes suprasegmental information relative to the (segmental) point of recognition in a given word.

Methods: Brazilian Portuguese speakers (n = 51) participated in an auditory lexical decision task with trisyllabic words (n = 390). Nonce words (n = 195) were created by changing the onset consonant of the final syllable of real words (n = 195)—the point of recognition. For example, from the real word *trópiço* 'tropic', the nonce word *trópiro* was generated. The stimuli were controlled for onset cluster quality and position, coda quality and position and, crucially, stress. As well, lexical frequency and neighbourhood density were analysed for each real word, and phonotactic bigram probability was calculated for all words based on the Portuguese Stress Corpus (Garcia 2014). RTs were measured from the end of each stimulus. The data were modelled with Multilevel Linear Regression (`lmer()` in R) with by-speaker and by-item random intercepts. RTs were normalized by speaker, centred and scaled. Frequency, neighbourhood density and bigram probability were centred and scaled. Only participants whose accuracy level was above or equal to 80% (n = 37) were considered in the analysis.

Results: Consistent with the literature (e.g., Carreiras et al. 1997), real words had faster RTs ($\hat{\beta} = -0.07, p = 0.04$). Word frequency had an impact on RT, i.e., more frequent words yielded faster RTs ($\hat{\beta} = -0.05, p < 0.0001$). Neighbourhood density also had an effect ($\hat{\beta} = -0.02, p = 0.007$): words with more neighbours had faster RTs, contrary to what is traditionally found for English (e.g., Sommers 1996, Kirk et al. 1994, among others), but consistent with what Vitevitch and Rodríguez (2004) found for Spanish. Crucially, words with APU stress yielded *faster* RTs than words with PU ($\hat{\beta} = -0.04, p = 0.0006$), which yielded faster RTs than words with U stress ($\hat{\beta} = -0.06, p < 0.0001$). These results suggest that speakers are slower when the segmental point of recognition and suprasegmental information must be processed at (nearly) the same time. Thus, words with final stress have longer RTs. The fact that words with the most common stress pattern in Portuguese (PU) do not have faster RTs than words with APU stress indicates that, in the case of English, the stress effects on RTs are likely not due to a frequency bias toward word-initial stress.

Sonorant-conditioned mid vowel lowering in Turkish

Deepthi Gopal, University of Manchester
deepthi.gopal@manchester.ac.uk

Stephen Nichols, University of Manchester
stephen.nichols@manchester.ac.uk

In Turkish (Göksel & Kerslake 2005:10), /e/ lowers to [æ] in syllables with coda {m, n, r, l}, thus, /sen/ [sæn] 'you', /termɑl/ [tær.mal] 'thermal', /gizem/ [gi.zæm] 'mystery', /gelmek/ [gæl.mek] 'to come'. As far as we are aware, this pattern remains unanalysed in phonological terms. In this work, we offer new acoustic evidence that this lowering is *categorical*, provide a perceptual and phonetic analysis, and evaluate targeting of other mid vowels (/ø/ and /o/). We base our discussion on data from 5 female and 2 male Turkish speakers, geographically diverse but controlled for age (31–38, mean 34; from İstanbul, Bursa, Fethiye, Ankara (2), Denizli and Kars). Speakers read a list of 190 items in isolation and a further 35 sentences containing tokens of /e/ embedded in varied phonological and morphological environments; F1, F2 (at the mid-point of the vowel), and duration (from first to last visible period) were measured.

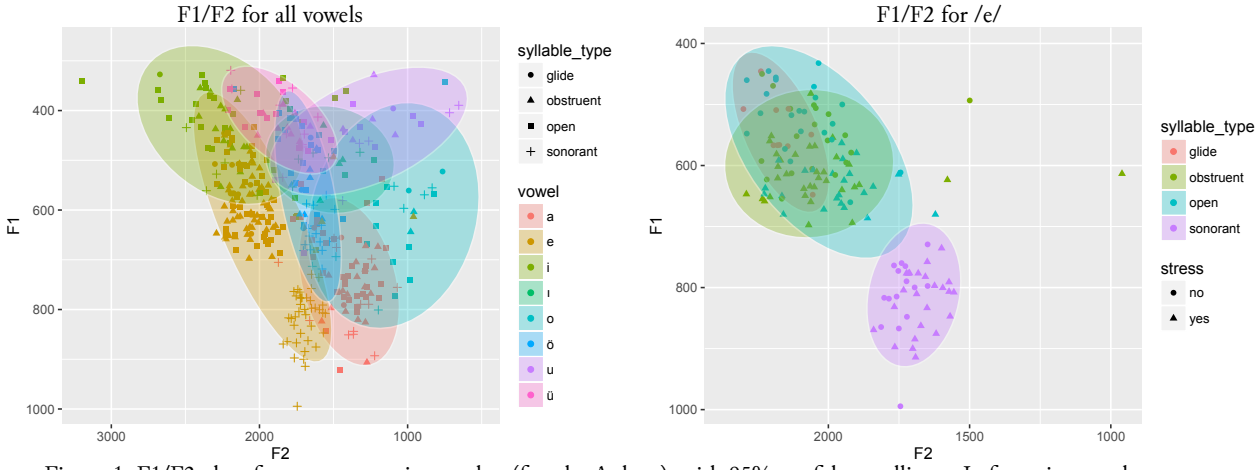

Figure 1. F1/F2 plots for a representative speaker (female, Ankara), with 95% confidence ellipses. Left: entire vowel space, vowels marked orthographically (note *ö* /ø/-lowering with sonorant codas, in blue). Right: /e/ alone for the same speaker, showing nonoverlapping clusters for sonorant and non-sonorant codas.

Our results confirm that speakers have two *categorically* distinct realisations of /e/ conditioned by the following coda, except in the dialects of far Eastern Turkey (Kars) in which there is some (gradient) lowering, especially before /r/. Lowering in *sonorant*-coda syllables is duration-independent and stress-independent, i.e. we find no statistically significant differences between unstressed and stressed syllables; stressed open syllables display noticeable gradient lowering relative to unstressed ones. There is inter-speaker variation for the other mid vowels: no speakers show predictable lowering in /o/, but /ø/ is systematically lowered to [œ] before sonorant codas by western (İstanbul, Bursa, Fethiye, Denizli) speakers and one central speaker (Ankara). There is slight overlap between lowered /ø/ in sonorant-coda syllables and lowered /ø/ in final open syllables, which is not seen for /e/, and not all /e/-lowering speakers also apply this rule to /ø/.

Although cases of lowering in closed syllables are known (cf. Féry 2003 on closed-syllable laxing in French, also Storme 2014), the salience of *sonorancy* requires an analysis not predicated on duration, stress, or mora, as many previous analyses are. We suggest that lowering in sonorant-coda syllables is driven by various perceptual requirements: to improve the perceptibility of vowel-sonorant transitions and manner contrasts in sonorants, and to compensate for perceptual lowering driven by anticipatory nasalisation. Briefly, we note that vowel-sonorant transitions are poorly cued relative to vowel-obstruent transitions (cf. Fleischhacker 2002); F1 is depressed for both nasals and laterals (Johnson 2003), suggesting that F1-raising in the preceding vowel will enhance contrast. In an (anticipatorily) nasalised vowel (ibid.), due to bandwidth widening of the existing formants and the presence of antiformants, speakers tend to misperceive F1 as raised, which then drives vowel lowering (with the misperception feeding the production of raised F1) – this may then generalise to other [+sonorant] segments. An ongoing follow-up study will confirm the extent to which this is reflected in coda discrimination tasks.

High vowel deletion in Québec French: Evidence for vestigial iambs

Natália Brambatti Guzzo, Heather Goad, Guilherme Duarte Garcia (McGill University)
nataliaguzzo@me.com, heather.goad@mcgill.ca, guilherme.garcia@mail.mcgill.ca

Research in prosodic phonology has shown that languages organize segmental material into a limited number of prosodic constituents, as in (1) (Nespor & Vogel 1986, McCarthy & Prince 1995).

(1) PPh
|
PWd
|
Ft
|
σ

Phonological processes are identifiers of particular prosodic domains: stress, for example, is computed in the phonological word (PWd) and realized in the foot (Ft), as exemplified for English *avocado*: $[(\text{ˌævə})_{Ft}(\text{ˈkɑː})_{Ft}\text{doʊ}]_{PWd}$. Although the constituents in (1) are typically assumed to be universal (e.g. Selkirk 1996), the validity of this assumption, notably for the foot, has been challenged by research on French: the only position of obligatory 'stress' is the right edge of the phonological phrase (PPh), independent of the number of lexical words it contains: [lə mɔvɛ avɔˈka] 'the bad avocado' (Dell 1984). This has led some researchers to analyse French 'stress' as intonational prominence and French, unlike the vast majority of languages, as a language without the foot (e.g. Jun & Fougeron 2000).

In this paper, we examine a different process, high vowel deletion (HVD), in Québec French (QF) and argue that it motivates *iterative iambic footing* in this variety, even though typical signatures of word level stress are absent. We focus on [i], the most deletable vowel (after [ə]) in QF.

Native speakers of QF (*n*=5) rated how natural 2-6 syllable words (*n*=355) sounded with 0-2 [i] vowels deleted in various non-final positions in the word. Representative examples are in (2).

(2)	HVD site (underlined)		Position(s) of HVD from right edge
a.	[kɔ̃.bi̲.ˈne]	'to combine'	Syllable 2
	[ma.ni̲.fɛs.ta.ˈsjɔ̃]	'demonstration'	Syllable 4
	[re.si̲.prɔ.si̲.ˈte]	'reciprocity'	Syllables 2 & 4
b.	[ɔr.ga.ni̲.za.ˈtœr]	'organizer'	Syllable 3
	[y.ni̲.vɛr.sa.li.ˈte]	'universality'	Syllable 5
	[ka.pi̲.ta.li̲.za.ˈsjɔ̃]	'capitalization'	Syllables 3 & 5

The data were modelled with a multilevel ordinal regression (by-speaker and by-item random intercepts). The first result is that medial HVD is preferred over initial HVD ([kɔ̃.bi̲.ˈne] > [fi̲.ˈlɛ] 'net') ($\hat{\beta} = 0.93, p < 0.05$). We hypothesize that initial HVD is dispreferred as the first syllable of the leftmost PWd in a PPh typically bears high tone in French and, thus, is prominent (Jun & Fougeron 2000, Thibault & Ouellet 1996 on QF); initial deletion may also impede lexical access.

Second, there is an interaction between HVD in initial vs. medial position and HVD in even (2a) vs. odd (2b) numbered syllables from the right edge ($\hat{\beta} = 1.13, p < 0.05$): in medial position, deletion of vowels in even numbered syllables is preferred (Verluyten 1982; cf. Cedergren 1986). If QF builds iambic feet iteratively from right-to-left, HVD in type (2a) words will be preferred, as these vowels are located in foot-dependent position, as shown in (3). Heads of feet are in bold.

(3)				
	HVD preferred (2a):	[kɔ̃(bi̲.ˈ**ne**)]	[ma(ni̲.**fɛs**)(ta.ˈ**sjɔ̃**)]	[re(si̲.**prɔ**)(si̲.ˈ**te**)]
	HVD dispreferred (2b):	[ɔr(ga.**ni̲**)(za.ˈ**tœr**)]	[(y.**ni̲**)(vɛr.**sa**)(li.ˈ**te**)]	[(ka.**pi̲**)(ta.**li̲**)(za.ˈ**sjɔ̃**)]

Third, HVD yielding strings mirroring licit complex onsets is dispreferred ($\hat{\beta} = -0.90, p < 0.05$): deletion in [bi̲.ˈʒu] 'piece of jewellery' and [kɔ̃.bi̲.ˈne] is preferred over deletion in [fi̲.ˈlɛ] and [su.pi̲.ˈre] 'to sigh' as [bʒ], [bn] are illicit onsets, while [fl], [pr] are licit. We propose that this shows that syllabification and footing remain intact and that HVD affects segments only, e.g. [kɔ̃.bØ.ˈne]. That deletion is not equally preferred in [kɔ̃.bi̲.ˈne] and [su.pi̲.ˈre] we attribute to recoverability. It is easier to recover the deletion site when the string of consonants resulting from HVD is phonotactically ill-formed: [kɔ̃bne] can only be reconstructed as [kɔ̃bVne], but [supre] could be [supre] or [supVre]. If HVD instead led to resyllabification and refooting, deletion would wrongly be favoured in cases where the resulting consonant string is phonotactically licit: [supre] > [kɔ̃bne].

Phonetics of Closed Syllable Shortening in Palestinian Arabic

Nancy Hall, California State University Long Beach

This acoustic phonetic study compares lexically short vowels in Palestinian Arabic to vowels that are underlyingly long, but have undergone 'closed syllable shortening' (Abu-Salim 1986, Younes 1995). The question is whether such phonological alternations fully or only partially neutralize underlying distinctions on a phonetic level.

Shortening was examined in two morphological environments, both of which involve repairing CV:CC syllables. The first environment occurs when the negative suffix /ʃ/ is attached to a CV:C root, as in /fa:q-ʃ/ 'woke-negative'. Post-shortening, the resultant CVCC syllable is further subject to optional epenthesis, producing [faqeʃ]. Note that epenthesis counter-bleeds shortening; the CC coda that triggered shortening is not present on the surface in [faqeʃ], although a transparent variant [faqʃ] exists as well.

The second environment occurs when the dative suffix /l/ is applied to a CV:C root, as in /ʒa:b-l-ak/ → [ʒablak] 'brought to you'. Here, the traditional analysis crucially relies on the phonological cycle: shortening applies to the intermediate form /ʒa:b-l/ 'brought to', before the addition of /-ak/ 'you' causes the /l/ to resyllabify as an onset. In this case, the CC coda that triggered shortening is never present at the surface, so the environment is fully opaque.

In an exploratory study of a small number of items produced in 8 repetitions by 72 speakers, shortened vowels were found to be statistically indistinguishable in duration from underlying short vowels. The first syllable vowel of [ʒablak] (/ʒa:b-l-ak/) had the same duration as that of [ʒabha] (/ʒabha/) 'front'. The vowels of words like [(ma)faqeʃ] (/(ma)fa:q-ʃ/) had the same durations as those of words like [naqeʃ] (/naqʃ/) 'engraving'.

This finding contrasts with results for another process, vowel epenthesis, which did show incomplete neutralization in the same experiment. I suggest that complete phonetic neutralization may be more common in processes whose conditioning environment is opaque, like shortening, than in more transparent processes.

It was also found that speakers differ as to the contexts in which they apply shortening: speakers from some regions did not shorten vowels before the dative suffix /-l/ (that is, they produced [ʒa:blak]), but did shorten vowels before the negative suffix /ʃ/. This suggests that closed syllable shortening in these two morphological environments is not a unified phenomenon, as it has sometimes been analyzed.

Mora Preservation in Gújjolaay Eegimaa Reduplication

Abbie Hantgan (SOAS), Serge Sagna (University of Surrey), Stuart Davis (Indiana University)
ah104@soas.ac.uk, eakkut@gmail.com, davis@indiana.edu

Gújjolaay Eegimaa (GE) is a variety of Jóola spoken in Senegal. The most well-known variety of this language is Diola-Fogny from Sapir (1965), discussed by Ito (1986) and Kager (1999) among others. In both these varieties, a word-internal coda respects the Coda Condition (CodaCon), as in Ito (1986), where the coda doesn't license place features and surfaces as either a nasal homorganic to the following obstruent onset or as the first part of a geminate. Also, both varieties witness consonant deletion as a repair for input sequences whose faithful output violates CodaCon, and both permit word-final codas. GE differs from other varieties in lacking phonemic long vowels, but permits vowel hiatus. With this as background, an intriguing aspect of GE is its reduplication pattern as in the perfective, which entails the (prefixal) reduplication of the verb root (mainly monosyllables). Data in phonemic transcription are in (1) and (2) (ni =1st pers. sg.).

	Root	Faithful Redup.	Output	Gloss		Root	Faithful Redup.	Output	Gloss
1a	/cɔb/	/nɪ-cɔb-cɔb/	*nɪ-cɔccɔb*	'frighten'	1b	/bɛr/	/nɪ-bɛr-bɛr/	*nɪ-bɛbbɛr*	'laugh'
2a	/bɔk/	/nɪ-bɔk-bɔk/	*nɪ-bɔbɔk*	'dance'	2b	/pap/	/nɪ-pap-pap/	*nɪ-papap*	'dust'

(1) shows that if the reduplicant ends in a voiced consonant it totally assimilates to the following root-initial consonant whereas (2) shows that a voiceless consonant at the end of a reduplicant deletes. Both the assimilation strategy in (1) and the deletion strategy in (2) can be seen as a means to respect the CodaCon (as in Kager 1999 for Diola Fogny). If the faithful reduplicant would respect CodaCon then no relevant change occurs as in the reduplication of /gan/ 'lose weight' as *nɪ-gaŋgan*. While two previous analyses of GE reduplication (Sagna 2008, Bassène 2012) consider the difference between (1) and (2) being that voiced consonants assimilate while voiceless ones do not, we propose based on a broader consideration of the data that the total assimilation in (1) is motivated by mora preservation of the deleted moraic coda of the faithful reduplicant. Specifically in (1), from a derivational perspective, when the voiced coda consonant of the reduplicant deletes to satisfy CodaCon, gemination of the onset is triggered by mora preservation of the deleted voiced coda consonant. There is no gemination in (2) because single voiceless consonants are not assigned a mora (a sonority-based difference). Evidence for this conception comes from the reduplication of roots that begin with a vowel and from those that end in a geminate. First, consider roots of the structure VC in (3) and (4): when these reduplicate their outcome depends on the voicing of the root-final consonant.

	Root	Faithful Redup.	Output	Gloss		Root	Faithful Redup.	Output	Gloss
3a	/ap/	/nɪ-ap-ap/	*nɪ-a-ap*	'forge'	3b	/ɛs/	/nɪ-ɛs-ɛs/	*nɪ-ɛ-ɛs*	'slice'
4a	/al/	/nɪ-al-al/	*nɪ-awwal*	'ripe'	4b	/ɪb/	/nɪ-ɪb-ɪb/	*nɪ-ɪjjɪb*	'cut'

In (4) the deletion of a voiced consonant in coda position of the faithful reduplicant triggers insertion and gemination of a homorganic glide, whereas the deletion of a voiceless consonant results in vowel hiatus. The gemination in (4) can't be the result of assimilation, but is consistent with it resulting from mora preservation of the deleted coda; an output like [nɪ-al.lal] for (4a) would violate CodaCon since the features of the lateral [l] originate in coda position. There is no gemination in (3) since the deleted voiceless coda is not moraic. Second, while GE geminates are phonemic for all consonant types, only voiceless geminates occur word-finally. Consider the reduplicated form of /sɔpp/ 'hurt'. This reduplicates as *nɪ-sɔssɔpp*, where the deleted *geminate-p* triggers gemination of the root-initial onset whereas in (2b) the deleted coda singleton /p/ fails to trigger gemination. Such is consistent with a moraic view of geminates (Hayes 1989). Thus GE reduplication has implications for various aspects of moraic theory, the nature of CodaCon and the role of intermediate forms for an OT analysis in harmonic serialism or sympathy theory.

Phonotactics with[awt] rules: the learnability of a simple, unnatural pattern in English

John Harris (UCL), Nick Neasom (UCL) and Kevin Tang (Yale)

There is a well-established collection of speaker-independent methods for discovering phonotactic patterns in languages, e.g. comparative reconstruction, phonological analysis and computational learning. There is also an increasingly varied collection of experimental methods for ascertaining how much of this patterning is actually internalised by speaker-hearers. In seeking to determine what makes a phonotactic pattern learnable or not, researchers have focused on a variety of factors, including phonological regularity, productivity, naturalness, and formal simplicity. Experimental studies have investigated various permutations of these factors, with results that are more or less surprising. For example, speakers have been shown to have internalised and to be able to productively apply (a) patterns that are regular, simple and natural (e.g. wug tests of English **-s**) but also (b) patterns that are irregular, relatively complex and not synchronically natural, such as English velar softening (e.g. Pierrehumbert 2006).

In this paper, we examine the English phonotactic pattern where consonants following /aw/ are restricted to coronals; hence **tout**, but not */tawk/, */tawp/ (e.g. Halle & Clements 1983). The pattern ('awT') is pretty regular, more so than velar softening. It is general, in that it affects a large swath of the lexicon. It is formally quite simple, more so than the **-s** pattern. And it is not natural. It is the synchronically accidental outcome of a series of largely unrelated sound changes; each of the changes might be natural, but their cumulative effect is not.

We report the results of two non-word judgement experiments designed to test the extent to which native speakers of English have tacit knowledge of the awT pattern. In both experiments, listeners were presented with $C_{1\text{-}3}$VC non-word stimuli containing the diphthongs /aw/, /ow/, /ij/, followed by a range of consonants, and were asked to judge how English-like they sounded. The non-words were controlled for lexical neighbourhood density, weighted by frequency. In the first test, listeners made forced choices between paired words distinguished solely by whether the vowel was followed by a coronal versus a non-coronal consonant. In the second, listeners rated individually presented stimuli on a scale of Englishness, drawn from a sample of ~1200 nonwords.

The question of whether speakers have implicit knowledge of a given phonotactic pattern can be approached in two stages: (a) do they have any tacit awareness of the pattern at all and, if so, (b) is the awareness commensurate with the pattern being stored as a grammatical rule? Broadly speaking, the results of both experiments show weak evidence of an awareness of awT but little or no evidence that this reflects grammaticalised knowledge. That is, to the extent that speakers have any tacit inkling of the pattern at all, it is probably not encapsulated in anything like a phonologist's rule or constraint. Where a coronal preference is observable, it does not generalise across different manners of articulation, as would be expected if there were a rule-driven bias towards formal simplicity. Also, the preference is influenced by onset size and lexical neighbourhood factors, which suggests subjects were making on-the-fly judgements of how much the non-words resemble real words.

We conclude that awT is a case where phonologists know more about a phonotactic pattern than speakers know. In the light of our results, we consider whether this should be attributed to the fact that awT is not natural (cf. Hayes & White 2013) or to other factors, such as that it is not involved in alternations.

Typological predictions of an interactive learning model

Coral Hughto - UMass Amherst (coralwilliam@linguist.umass.edu)

Introduction: Combining a grammatical theory with a learning model generates probabilistic typological predictions over the range of possible grammars (see Staubs 2014). The work presented here focuses on the consequences of pairing an interactive agent-based learning model with a Maximum Entropy (MaxEnt) grammar framework (Goldwater & Johnson 2003). This pairing has the effect of increasing the predicted typological frequency of patterns which are easier to learn - that is, which can be generated with a greater range of possible constraint weights - and decreasing the predicted frequency of harder patterns.

A major consequence of this approach is that gang effects are predicted to be relatively rare, typologically. Gang effects occur when multiple violations of lower-weighted constraints outweigh a single violation of a higher-weighted constraint, and have been cited as a problem for weighted-constraint grammars because they allow for the representation of some unusual patterns (see Legendre et al. 2006). The interactive learning model (ILM) predicts that gang effects should be rare because these patterns require more precise relationships between constraint weights, which makes them harder to learn.

Case Study: This prediction is illustrated with the case study of the typology of contrasts between /s/ and /ʃ/: No Contrast (NC), Full Contrast (FC), Complementary Distribution (CD), Contextual Neutralization (CN; e.g. /s/ → [ʃ] before front vowels), and Elsewhere Neutralization (EN; e.g. /ʃ/ → [s] before non-front vowels). EN is not representable in parallel OT with standard constraints, but is representable in a weighted-constraint grammar as a gang effect. This pattern is attested in only one language, Gujarati (Carroll 2012).

Interactive Learning: The ILM generates typological predictions by calculating the percentage of runs in which the agents learned each possible pattern out of the total number of runs. Learning consists of the interaction between two agents taking turns in the roles of "teacher" and "learner". At each learning step, an input is selected at random, and each agent selects an output according to its current grammar. If the learner's form differs from the teacher's, the learner updates its constraint weights via gradient ascent. This process is repeated until there is at least 95% probability on each optimal output.

Results: The distribution of languages resulting from 1,000 runs of this simulation is reported in the table. The r^2 values show that the predictions of the ILM provide a better fit to the observed typological frequencies (reported in Carroll 2012) than frequency predictions derived by sampling the constraint weights 1,000 times. Crucially, the ILM predicts the relative rarity of the gang effect, while sampling overgenerates.

Lang.	Observed	ILM	Sampling
NC	44%	46.6%	13.7%
FC	37%	48%	40.7%
CD	10.3%	2.6%	6.6%
CN	8.2%	2.7%	5.6%
EN	0.5%	0.1%	32.2%
r^2		0.96	0.05

Discussion: The results from the ILM simulations illustrate the prediction that gang effect patterns should be typologically rare, and show that it mirrors a real typological trend. Future work will test whether this result extends to larger problems and other typologies. Overall, this work points to an asymmetry between representational capacity and typological prediction - the ability to represent a pattern does not imply the prediction that it should be well attested in typology - and highlights the influence of learning on typological trends.

The role of empty-headed syllables in subtractive morphology

Björn Köhnlein, Ohio State University, koehnlein.3@osu.edu

The issue. I argue that some cases of subtractive morphology can best be analyzed as an epiphenomenon of prosodic affixation that leads to word-final empty-headed syllables. This is in line with the idea that morphology is generally additive, and that non-concatenative morphology can be derived from prosodic / featural affixation (cf. Bermúdez-Otero 2012); yet it relativizes recent proposals which seem to suggest that defective integration of mora affixes into syllables with (vocalic) heads is sufficient to analyze subtraction (e.g. Trommer & Zimmermann 2014). Evidence for our claims comes from various German dialects that delete some word-final obstruents in plural forms. The precise contexts differ among dialects; for purposes of illustration, I focus here on the sequence /nd/. A representative example from the *Mittelrheinischer Sprachatlas (MRhSA)* (2002) is provided in (1):

(1) /hʊnd/ [hʊnt] 'dog' vs. /hʊnd/ + /'?'/ [hʊn] 'dogs' (MRhSA, Nussbach dialect)

The traditional analytical approach to such cases of subtractive pluralization is constraint-based, essentially requiring plural forms to end in a sonorant, which causes deletion, but not singular forms, where no deletion occurs (Golston & Wiese 1995, Knaus 2003). This analysis successfully accounts for the singular-plural alternations, but it misses an important generalization: as exemplified in (2), deletion in relevant dialects also occurs (predictably) in word-medial onsets of unstressed syllables (/ndV/ → [nV]; see, e.g., Schirmunski 1962).

(2) /kɪnd/ [kint] 'child' vs. /kɪnd/ + /ɐ/ [kɪnɐ] 'children' (MRhSA, Nussbach dialect)

Proposal. In a nutshell, I argue that word-final deletion in (1) is a subset of 'weak onset deletion', which derives from the affixation of a trochaic foot template ($\sigma_s.\sigma_w$). In singulars, default syllabification puts word-final /d/ in coda position (final devoicing applies) – [hʊnt]. In plurals, the morpheme ($\sigma_s.\sigma_w$) would force /d/ into the onset of a weak (empty-headed) syllable, which leads to deletion. (Roughly speaking, the more sonorous nasal makes a better onset of the post-tonic syllable.) The outcome is [hʊ.n] (dot = syllable boundary). The analysis works in the same way for 'overtly' disyllabic plurals, such as [kɪnt] vs. [kɪ.nɐ].

Particularly strong evidence in favor of the approach comes from Franconian tone accent dialects, many of which display alternations of the type in (1) and (2). These dialects also have tonal alternations between Accent 1 and Accent 2, which some have analyzed as a contrast between monosyllabic (Accent 2) and disyllabic feet (Accent 1, sometimes with an empty-headed second syllable); see, e.g., Köhnlein 2011, to appear, Van Oostendorp to appear. Crucially, subtractive pluralization seems to *always* coincide with a change from Accent 2 to Accent 1, which is consistent with the (independently motivated) claim that Accent 1 is disyllabic. For instance, the map 'dog / dogs' from the MRhSA features 252 accent dialects with subtraction (as well as 196 non-accent dialects); in all 252 accent dialects, consonant deletion is accompanied by a switch from Accent 2 to Accent 1.

Theoretical implications. The analogous behavior of word-final and word-medial consonants seems hard to capture with an approach where associating suffix moras to word-final coda consonants is regarded as the sole trigger of such subtractions. Word-final deletion in non-tonal dialects might be analyzed along these lines (as suggested in Trommer & Zimmermann 2015: 489), but the analysis does not capture word-medial onset deletion (as I shall discuss, however, mora affixation along the lines of van Oostendorp 2006, which can also lead to empty-headed syllables, might offer an empirically adequate alternative to the foot-based analysis). Furthermore, it seems non-trivial to integrate the Franconian tonal facts into a moraic approach. In sum, our talk provides further confirmation that prosodic affixation is a promising approach to subtractive morphology and illuminates the role that empty-headed syllables / foot structure can play in shaping relevant alternations.

Wrong side reduplication in Koasati

Leland Kusmer and Ivy Hauser (University of Massachusetts Amherst)
(lkusmer)|(ihauser)@linguist.umass.edu

Introduction: Marantz's frequently cited generalization (1982) states that reduplicant material tends to be adjacent to the corresponding base material. Nelson (2002) has claimed that observed cases of "wrong side reduplication" (WSR) are actually epiphenomenal and thus not true counterexamples to Marantz's generalization. We argue that Koasati (Muskogean, Kimball (1986)) verbal reduplication constitutes true WSR which supports models that can generate both wrong and adjacent side reduplication. We propose a model based on positional faithfulness, making the prediction that WSR should only copy material from strong positions.

Wrong side reduplication: Nelson (2002) examined apparent counterexamples to Marantz's generalization, and proposes that all instances can be explained as other phenomena and not true WSR. For some cases she proposes a full copy plus deletion analysis; for others, a non-reduplicative copying process to meet a prosodic template. Most grammatical models can generate both adjacent reduplication and WSR, which is a problem if no true cases of WSR exist.

Koasati data: Koasati has a reduplication process used to form event plurals. The initial consonant of the first CV syllable + [o] is suffixed to the stem.

	Stem	Reduplicated	Gloss
(1)	cofok-	cofok-**co:**-	be angled
(2)	lapat-	lapat-**lo:**-	be barren
(3)	limih-	limih-**lo:**-	be smooth
(4)	alot-	alot-**lo:**-	be full
(5)	pak-	pak-**po:**-	have a blister

The Koasati case cannot be analyzed as full copying with deletion as there is no independently attested process in the language that would delete an entire stem except the first consonant. Even if we were to propose such a process, it would need to delete material from both sides as in (4), which would be a highly unusual deletion process. Koasati WSR also cannot be non-reduplicative copying as there is no single prosodic template the reduplicated forms meet. For example, it is not the case that the reduplicated stems must always be trisyllabic, as evidenced by (5).

Riggle (2004) reports that Creek, another Muskogean language, also has a true WSR process involving right aligned infixes which copy from the left of the root. In contrast, the Koasati pattern is fully suffixing, providing a clearer example of this phenomenon and strengthening the argument that WSR must be included in our typology.

Analysis: In order to account for this process, the analysis must copy material from the left side while suffixing that material. The material then undergoes melodic overwriting to meet the C[o:] template. We provide an analysis which relies on positional faithfulness (Beckman 1998) to the left edge of the root to fill the suffixed reduplicative template. Constructing the base on the surface via root-base faithfulness (Downing 1998) accounts for the base shrinkage in vowel initial words like (4). This analysis will permit limited WSR in the cases of suffixes copying root-initial material, maintaining most of Marantz's generalization while allowing for cases like Koasati.

Conclusions: We have identified a new example of WSR, showing that true WSR must be included in our grammatical models. The data here cannot be explained by other non-reduplicative processes. However, the copying is limited to strong positions such as the initial consonant of the root, which have been independently argued to require special faithfulness. Therefore we predict WSR to exist, but only in cases like Koasati where the reduplicated material is copied from the left edge but aligned to the right edge.

The Status of Underlying R in Haitian: representational solutions to counterfeeding

Mohamed Lahrouchi
CNRS – Université Paris 8

Haitian, a French-lexicon based Creole (HC), has a phonemic nasal vowel contrast (e.g. [bɑ̃k] 'bank' / [bak] 'tub'). It has also developed a regressive optional nasal harmony whereby an oral vowel is nasalized when immediately followed by a nasal consonant (e.g. [fɑ̃mi] 'family'). Any segment intervening between the vowel and the nasal consonant blocks nasalization: /kalm/ *[kɑ̃lm] 'calm'. Curiously, this harmony interacts with another process: coda R-deletion ([tɛ] *terre* 'ground' (with no accompanied compensatory lengthening cf. [kɑ:] 'car')). Consequently, in accordance with nasalization, *chambre* 'room' surfaces as [ʃɑ̃m], while *charme* 'charm' surfaces as [ʃam] without either R or nasalization where we could have expected: [ʃɑ̃m]. This constitutes a classical case of phonological opacity, more examples are shown in (1) beneath.

(1)	a. Nasal Harmony (NH)		b. R-deletion		c. R-deletion blocking NH	
	lɑ̃m	'blade'	efɔ	'effort'	lam	'tear'
	pɑ̃n	'breakdown'	libɛte	'freedom'	bɔn	'landmark'

The examples in (1a) illustrate regressive nasal harmony, while those in (1c) show cases where r-deletion in post-vocalic position counter-feeds nasal harmony. Forms in (1b) show contexts where /r/ does not surface. Ans (1968), Cadely (1994) and Nikiema & Bhatt (2005), among others, argued that the forms that resist nasalisation have an underlying /r/, while only a few refute this hypothesis Valdman (1978). Non-derivational theories face the problem of how to explain that the same context results in different surface forms. How to account for this blocking of harmony without having to propose an opaque counter-feeding rule ordering, and without having to posit an underlying R in forms where it is *never* visible: V[R]N.

Our solution is the following, in Haitian R must be licensed by being in a pre-vocalic environment. In the creolisation, any R in a pre-consonantal context was unlicensed, but what followed was merely the deletion of the segment (R), not its syllabic structure (CV). Following Scheer (2012), in V[R]N forms such as [lam] 'tear', we propose that the remnant of R, an empty CV, disrupts the locality between the target of nasalization (a vowel) and its trigger (a nasal): [la-CV-m]. This empty CV blocks harmony just as the /l/ does in [kalm] 'calm' (as shown in 2). The implication of this analysis is a representational solution to the counterfeeding analysis (deletion blocks harmony).

However, this is not to say that there is no underlying R in Haitian. Against the previous literature, underlying R is only posited where there are alternations: [tɛ] 'ground' and [libɛte] 'freedom' vs. [ɑ̃tere] 'bury' and [libere] 'to free'. These alternating forms include a floating R. However, nouns ending in R (/ter/) behave exactly like vowel-final nouns (cf. Nikiema & Bhatt 2005 and Nikiema 1999 for the details). To complete the analysis, we explain how R selects vowel-seeking allomorphy. We propose a derivation by phase account with its accompanied phase impenetrability condition (Chomsky 2001, Marvin 2002 and Marantz 2007, among others). Nouns like /ter/ are spelled-out within nP, at this point the R loses its chance to be licensed. Consequently, the final /r/ is not visible when the definite article is added, resulting in a vocalic allomorph -a in stems that are ended by floating R (as shown in 3).

(2a)
```
C  V  C  V  C  V
|  |        |
ʃ  a        m      > [ʃam]
k  a  l     m
```

(2b)
```
C  V  C  V
|  |  |
ʃ  a  m        > [ʃɑ̃m]
f  a  m  i     > [fɑ̃mi]
```

(3) $[[ter]_{nP}\ la]_{DP}$ = 1. Spell out nP > [tɛ]
2. Hiatus /tɛ+la/ > [tɛa]

Opacity, output-drivenness, and metrical faithfulness constraints

Giorgio Magri | SFL (CNRS and University of Paris 8)

□ Traditionally, *opacity* is defined in rule-based terms (Baković 2011). How can opaque processes be singled out in constraint-based terms? Tesar (2013) calls a grammar G *output-driven* provided it satisfies (1): if G maps some UR $\mathbf{a}$ to some SR $\mathbf{d}$, it also maps to $\mathbf{d}$ any UR $\mathbf{b}$ which is more similar to $\mathbf{d}$ than $\mathbf{a}$ is, namely $(\mathbf{a},\mathbf{d}) \leq_{\text{sim}} (\mathbf{b},\mathbf{d})$ where $\leq_{\text{sim}}$ is a partial order among pairs of a UR and a SR based on their degree of *similarity*. To illustrate, the *saltation* pattern /p/→[β], /b/→[b] (White 2013) is non-output-driven because (/p/, [β]) $\leq_{\text{sim}}$ (/b/, [β]) and thus (1) fails with the positions $\mathbf{a}$ = /p/, $\mathbf{b}$ = /b/ and $\mathbf{d}$ = /β/. Saltations are opaque patterns. Tesar submits that in general opaque patterns can be defined as the non-output-driven ones.

(1) **If:** $G(\mathbf{a}) = \mathbf{d}$ and $(\mathbf{a},\mathbf{d}) \leq_{\text{sim}} (\mathbf{b},\mathbf{d})$
Then: $G(\mathbf{b}) = \mathbf{d}$

□ Zooming in on OT, Tesar proves that OT grammars are output-driven provided each faithfulness constraint F satisfies the *faithfulness output-drivenness condition* (FODC$^{\text{OT}}$) in (2) for any candidates $(\mathbf{a},\mathbf{d}) \leq_{\text{sim}} (\mathbf{b},\mathbf{d})$, any other candidate $(\mathbf{b},\mathbf{c})$, and some candidate $(\mathbf{a},\mathbf{c})$. Does this FODC$^{\text{OT}}$ admit an intuitive interpretation?

(2) **If:** $F(\mathbf{a},\mathbf{d}) < F(\mathbf{a},\mathbf{c})$ **Then:** $F(\mathbf{b},\mathbf{d}) < F(\mathbf{b},\mathbf{c})$
If: $F(\mathbf{b},\mathbf{c}) < F(\mathbf{b},\mathbf{d})$ **Then:** $F(\mathbf{a},\mathbf{c}) < F(\mathbf{a},\mathbf{d})$

□ Intuitively, a faithfulness constraint F measures the distance between URs and SRs along some phonological dimension. A *distance* (or *metric*) is a function *dist* which pairs two points A and B with a number $dist(A,B)$ in compliance with certain axioms. A crucial axiom is the *triangle inequality*: the distance $dist(A,C)$ is never larger than the distance $dist(A,B)$ plus the distance $dist(B,C)$ for any choice of the point B, as in (3). These considerations motivate the *faithfulness triangle inequality* (FTI) in (4).

(3)

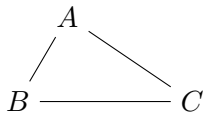

(4) $F(\mathbf{a},\mathbf{c}) \leq F(\mathbf{a},\mathbf{b}) + F(\mathbf{b},\mathbf{c})$

□ Tesar offers a concrete definition of the similarity condition $(\mathbf{a},\mathbf{d}) \leq_{\text{sim}} (\mathbf{b},\mathbf{d})$. I show that his definition satisfies the axiom (5): the less similar candidate $(\mathbf{a},\mathbf{d})$ "makes up" not only for any faithfulness violation incurred by the more similar candidate $(\mathbf{b},\mathbf{d})$ but also for any faithfulness violation incurred by the candidate $(\mathbf{a},\mathbf{b})$ which pits $\mathbf{a}$ and $\mathbf{b}$ against each other. Suppose that F satisfies the FTI (4). Suppose that $(\mathbf{a},\mathbf{d}) \leq_{\text{sim}} (\mathbf{b},\mathbf{d})$. Hence, (5) holds, whereby $F(\mathbf{a},\mathbf{b}) \leq F(\mathbf{a},\mathbf{d}) - F(\mathbf{b},\mathbf{d})$. Replacing the latter in the FTI (4), yields $F(\mathbf{a},\mathbf{c}) \leq F(\mathbf{a},\mathbf{d}) - F(\mathbf{b},\mathbf{d}) + F(\mathbf{b},\mathbf{c})$. Equivalently, $F(\mathbf{a},\mathbf{c}) - F(\mathbf{a},\mathbf{d}) \leq F(\mathbf{b},\mathbf{c}) - F(\mathbf{b},\mathbf{d})$. Thus if $F(\mathbf{a},\mathbf{c}) - F(\mathbf{a},\mathbf{d})$ is positive and the antecedent of the first implication in (2) holds, then $F(\mathbf{b},\mathbf{c}) - F(\mathbf{b},\mathbf{d})$ is positive as well and the consequent of the first implication in (2) holds. Analogous considerations hold for the second implication in (2). In conclusion, my **first contribution** is that the FTI entails the FODC$^{\text{OT}}$ for any constraint F.

(5) $F(\mathbf{a},\mathbf{d}) \geq F(\mathbf{b},\mathbf{d}) + F(\mathbf{a},\mathbf{b})$.

□ The reverse entailment fails in general. Yet, McCarthy (2003) conjectures that all phonologically relevant constraints are categorical. As my **second contribution**, I formalize the notion of *categoricity* for faithfulness constraints within McCarthy and Prince's (1995) *Correspondence Theory*. And I prove that the FODC$^{\text{OT}}$ entails the FTI for categorical faithfulness constraints. In other words, the FODC$^{\text{OT}}$ and the FTI are equivalent for categorical faithfulness constraints.

□ To strengthen this connection between output-drivenness and the FTI, I turn to HG. I repeat Tesar's reasoning which lead to (2), and conclude that the *HG faithfulness output-drivenness condition* (FODC$^{\text{HG}}$) in (6) ensures output-drivenness of HG grammars. This FODC$^{\text{HG}}$ in (6) is stronger than the FODC$^{\text{OT}}$ in (2) (indeed, (2) can be shown to be equivalent to (6) with ξ limited to $-1 < \xi < +1$). This is expected: a stronger condition is needed to discipline to output-drivenness the possibly larger typology of HG grammars. My **third contribution** is that the FODC$^{\text{HG}}$ is strong enough to be equivalent to the FTI for any (also non categorial) F.

(6) For every threshold ξ:
If: $F\big(\mathbf{b},\mathbf{c}\big) \leq F\big(\mathbf{b},\mathbf{d}\big) + \xi$
Then: $F\big(\mathbf{a},\mathbf{c}\big) \leq F\big(\mathbf{a},\mathbf{d}\big) + \xi$

□ To conclude, understanding opacity in constraint-based phonology means understanding output-drivenness (1). This in turn means understanding the FODC$^{\text{OT}}$/FODC$^{\text{HG}}$ in (2)/(6). This paper contributes to this enterprise by establishing equivalences between the FODC$^{\text{OT}}$/FODC$^{\text{HG}}$ and the FTI (5). This equivalence holds for any faithfulness constraint in HG and for categorical ones in OT, thus distilling an implication of McCarthy's (2003) categoricity conjecture. This equivalence says that the FODC$^{\text{OT}}$/FODC$^{\text{HG}}$ can be interpreted as simply requiring faithfulness constraints to measure phonological distance "sensibly", namely in compliance with the triangle inequality. (Lack of) opacity in constraint-based phonology is tied to the metrical nature of faithfulness constraints.

Liaison and propagation between phonology and morpho-syntax

Vania Masutti — University of Padua
Giuseppina Silvestri — University of Cambridge (UK)

Aim: In this study we compare two phonological phenomena in Romance, i.e. liaison (LSN) in French and propagation (PPG) in the Italo-Romance variety of Orsomarso (Cosenza). Following the assumptions of the phonology syntax-interface (Chomsky-Halle 1968; Selkirk 1984 a.o.), we aim to describe and interpret some morpho-syntactic contexts of activation of LSN and PPG within the nominal domain (DP). Our purpose is to verify whether these phonological phenomena can be equally interpreted as the spell-out of the head of Num(eral)P, i.e. as the realisation of number morphology on the N(oun).

Data: The LSN data are partly taken from the PFC online corpus (Durand, Laks-Lyche 2009), and partly based on the judgments of native speakers of Ile de France, Lorraine, Alsace. As for PPG, the evidence has been collected through fieldwork investigations with informants from Orsomarso (Cosenza).

Description: French LSN consists in the pronunciation of a word-final consonant (normally silent) in *sandhi*, which also produces re-syllabification. The application of LSN is sensitive to both phonological (1a) and syntactic conditions (1b) (Selkirk 1984, Durand 1990). LSN is realised in both plural and in singular contexts. In plural contexts (2), it is obligatory between the D(eterminer) and N (2a) and between the numeral and N (2b).

(1) a. *deux enfants* > [døzɑ̃fɑ̃] 'two children'
b. *deux en maths* > [dø ɑ̃ mat] 'two in maths'

(2) a. *des employées* > 'some employers'
b. *trois oiseaux* > 'three birds'

However, as for numerals, LSN is obligatorily given between the numeral *cent* and N, if and only if the number refers to 'hundred(s)' bigger than one (3c-d). If 'hundred'=1, i.e. *one*-hundred, (3a-b) LSN is not obligatorily realised, even though the phonological context would allow it:

(3) a. cent // euros
one hundred euros
b. cent // images d'archives
one hundred images of archives
c. deux-cent**s eu**ros
two hundreds euros
d. il y a quatre cent**s ha**bitants [PFC, loc. 42arf1]
there are four hundreds inhabitants

PPG is the output of a type of vocalic assimilation occurring in word-internal and *sandhi* contexts (Tuttle 1985; Savoia 1987; Schirru 2013). In a Syllable1-Syllable2 sequence (where Syllable2 must carry primary or secondary stress), PPG is a left-to-right spread of (part of/) the features of the phonological segments /u/, /o/ and /i/. It results in a change of vowel quality or, more often, in the insertion of a semivowel (4a-b):

(4) a. [ʌə pwakkə] < *[lu pakku] 'the package'
b. [ʌə pjakkə] < *[li pakki] 'the packages' (Terelle-Frosinone; Schirru 2013:85)

In Orsomarsese, within plural DPs PPG is very infrequent, being usually triggered by /u/ which coincidentally is the morphological mark of definite [ru] and indefinite [nu] D. Nonetheless, in plural DPs with numerals ending with a trigger-vowel, PPG is still ruled out (5a-b). However, PPG triggered by the numeral '*one*-hundred' in both plural (6a) and singular (6b) DPs is obligatory.

(5) a. /'otto 'kani/ > ['ɔttə 'kanə]
'eight dogs'
b. /a 'tʃintu 'tavuli/ > ['tʃɪntə 'tavələ]
'one hundred tables'

(6) a. /'tʃintu 'lire/ > ['tʃɪntu 'lwɪrə]
'one hundred liras'
b. /a 'tʃintu 'lira/ > [a 'tʃɪntə 'lwɪra]
'the (coin of) one hundred liras'

Analysis. When activated, both LSN and PPG are consistent signals of Number feature, set on plural and singular, respectively. The fine-grained decomposition of NumP reveals that the patterns of realisation of LSN and PPG with the numeral '*x*-hundred(s)' are outputs of morpho-syntactic properties. As for LSN, -/z/ is the plural mark on the N (in line with Selkirk 1972, cf. Delfitto-Schroten 1991:177-8, Morin-Kaye 1982:320), whereas -/t/ is never interpreted as a morphological mark of plural. In Italo-Romance PPG robustly signals the Number on singular/mass nouns, whereas in plural contexts the numerals are sufficient to license the non-singular/non-mass value.

Computational and learnability properties of conflicting long-distance dependencies

Kevin McMullin
University of Ottawa
kevin.mcmullin@uottawa.ca

Gunnar Ólafur Hansson
University of British Columbia
gunnar.hansson@ubc.ca

While most phonotactic co-occurrence restrictions hold between string-adjacent segments, certain types of interactions hold across intervening material. Such patterns may be assimilatory or dissimilatory in nature. For example, in Bukusu, two liquids in a word must agree in their specification of [lateral], resulting in a pattern of liquid harmony (/r…l/ → [r…r]). By contrast, Georgian requires that two co-occurring liquids disagree (/r…r/ → [r…l]). With respect to surface phonotactics, each of these isolated patterns belongs to a class of formal languages known as Tier-based Strictly 2-Local (TSL_2; Heinz et al. 2011). TSL_2 patterns have been shown to be computationally learnable (Jardine and Heinz *to appear*), and a number of experimental studies suggest that patterns fitting this description are learnable in the laboratory (Finley 2011, 2012, McMullin and Hansson 2014, McMullin 2016).

This paper addresses empirical, theoretical and learnability challenges that arise when conflicting phonotactic demands are simultaneously imposed on the same set of segments. In particular, we consider combinations of non-adjacent consonant (liquid) agreement (*l…r, *r…l) and disagreement (*l…l, *r…r) that co-exist in a language, but hold at different levels of locality: across an intervening vowel (transvocalic contexts, L̲VL̲) and at greater distances (beyond-transvocalic contexts, e.g. L̲VCVL̲). This yields the logical combinations in (1)–(2):

(1) Short-range harmony, longer-range dissimilation:
lVl, rVr; *lVr, *rVl *lVCVl, *rVCVr; lVCVr, rVCVl

(2) Short-range dissimilation, longer-range harmony:
*lVl, *rVr; lVr, rVl lVCVl, rVCVr; *lVCVr, *rVCVl

Such patterns are unattested cross-linguistically (or at best extremely rare) and the phonotactics of each of (1)–(2) are computationally relatively complex, belonging to at least the Locally Testable subclass of regular languages. This region has been hypothesized to be inaccessible to human learners, and experimental studies have found that learners exposed to patterns of this level of complexity tend to acquire a simplified version instead, one which does fall within the TSL_2 region (Lai 2012, McMullin 2016). From this perspective, we would expect neither (1) nor (2) to be human-learnable.

However, current theoretical models of long-distance consonantal phonotactics predict that certain types of conflicting dependencies like these should be possible. For example, Bennett's (2013) proposal that all dissimilation results from the avoidance of surface correspondence (Hansson 2001, Rose & Walker 2004) predicts that (1) is a possible sound pattern, while (2) falls outside of the factorial typology and is therefore, by definition, not learnable.

We present experimental data suggesting that neither of the above predictions is correct. In an artificial language learning study, adult native English speakers were trained on a language with two suffixes ([-li], [-ru]) which triggered alternations in preceding stems whenever these contained a liquid (CVCV[l/r]V or CV[l/r]VCV). Preliminary results indicate, contrary to the predictions outlined above, that human learners are capable of detecting both types of phonotactic regularities shown in (1)–(2), despite their inherent complexity.

We argue that what makes situations like (1) or (2) special is that although the full pattern is non-TSL_2, it can be decomposed into two independent TSL_2 restrictions, one overriding the other. For example, (1) can be defined as a ban against *rr and *ll sequences on a liquid tier ({l, r}), which is overridden by a ban against *rl and *lr on a consonant tier (consisting of {l, r, p, t, s, m, …}). Algorithmic solutions to the problem of discovering such component restrictions and their override relations from positive data are yet to be developed. We note, however, that a solution to this learning problem will have wider applicability, extending to other instances where one well-formedness generalization overrides ('outranks') another, rendering the latter non-surface-true.

Rhythmic vs. demarcational stress in Mapudungun

Benjamin J. Molineaux -- The University of Edinburgh
benjamin.molineaux@ed.ac.uk

The degree to which the overall phonology of a language appears to 'care' about word stress varies widely, as pointed out by Hyman (2014):

> 'Languages such as Hungarian or Turkish … seem different because their metrical structure has little or no relevance outside the stress system itself. The contrast with English, whose phonology cares so much about stress, is quite striking.'(59)

Mapudungun (aka 'Araucanian'; isolate; Chile/Argentina) has long been described as having phonologically determined word-stress (cf. Echeverría & Contreras 1965; Zúñiga 2006; Sadowsky et al. 2013). However, as with Turkish and Hungarian, the stress system participates only marginally in the phonology of the language, triggering no major phonotactic asymmetries, nor alternations in the vowels' contrastive system (cf. Sadowsky et al. 2013). Furthermore, views of the basic parameter-settings for Mapudungun stress find little agreement.

(1) Purported iambic and trochaic stress in Mapudungun:

	Reference	Foot type	Quantity-sensitivity	Direction	End rule
a.	Echeverría & Contreras (1965)	Iamb	Syllabic	Left-to-right	Left
b.	Zúñiga (2006)	Trochee	Moraic	Right-to-left	Right

Although the trochaic reading (cf. 1b) is favoured by most contemporary Mapudungun specialists (cf. Salas 2006; Zúñiga 2006; Sadowsky et al. 2013), Echeverría & Contreras' data (1a) has long been the standard source for Mapudungun within typological studies (cf. Hyman 1977; Kager 1993; Hayes 1995; Gordon 2002; Tesar 2004). In line with de Lacy (2014) and Molineaux (2016) – who critique the evidence for the iambic analysis – my own data will, for the most part, support the a trochaic parsing. This said, both analyses will be shown to have overlooked the key role stress plays in highlighting the language's morphological boundaries. In ignoring this aspect of the system, both accounts have shoehorned Mapudungun into an exclusively structural, rhythm-based analyses of stress, where I will claim a more functional, or demarcational role is required. Specifically, I will show how stem-final stress regularly takes priority over the right-aligned word-level trochee, eliminating it altogether, as in (2b).

(2) Stem (s) and word (ω) stress in Mapudungun (Molineaux 2014: 188)

a. [[ɨ.ˈʈʂɨf.]$_s$-tu.-pu.-ke.ˈla-j.-m-i]$_\omega$
throw-REST-TRLOC-HAB-NEG-IND-2-S
'You don't usually throw x back here'

b. [[e.lu-ɲ.ˈma.]$_s$-fi-j.-m-i]$_\omega$
give-APPL-DIR.3SP-IND-2-S
'You give him/her/it x'

Based on original fieldwork material, my analysis will show that, while Mapudungun is undoubtedly a stress language, it is the morphology – rather than the phonology – that seems to 'care' about stress. Lack of stress-based segmental asymmetries, absence of culminativity in words with multiple stresses (see 2a), as well as the unambiguous marking of the stem-edge, conspire to create transparency in the morphological structure. This is particularly relevant for a language with abundant, agglutinating suffixation. While languages of the fusional spectrum (such as English), can afford for the morphology to be subordinate to prosodic well-formedness, highly agglutinating languages such as Mapudungun appear to avoid it, striving to retain clear decomposability of the morphology.

The overall argument of the paper has important implications for the roles of the phonological and morphological word in Mapudungun, as well as for the precise nature of 'rhythm', 'demarcation' and 'culminativity' as definitional traits of stress across the world's languages.

Backward languages: The case of French *Verlan* in OT

Sophie Moracchini (MIT) - *smoracc@mit.edu*

This paper offers an analysis of French *Verlan*, a member of the larger class of Backward Languages (cf. Bagemihl 1989). I argue that in *Verlan*, a certain interaction of faithfulness constraints creates the illusion of syllable reversal, when in fact the system requires only the change of one immediate precedence relation between adjacent segments. All other components of the system reduce to faithfulness constraints and general markedness constraints of French.

Open monosyllabic words:	*fou* ('crazy')	/f.u/ → [u.f]
Closed monosyllabic words:	*louche* ('shady')	/lu.ʃ/ → [ʃə.lu]
Disyllabic words:	*moto* ('motorcycle')	/mo.to/ → [to.mo]
Polysyllabic words:	*cigarette'* ('cigaret')	/si.ga.ʁɛt/ → [ʁɛt.si.ga], [ʁɛt.ga.si],[ga.ʁɛt.si]

Triggering Verlanization The basic constraint specific to *Verlan* is Anti-Immediate-Precedence (Anti-IP) which requires that any word and its verlanized version of differ in one immediate precedence relation (→). The faithfulness constraints Max(C) and Max(V) militate for the preservation of segments in the ludling output, so that the loss of one immediate precedence relation automatically leads to the creation of a new one at the edge of the word, creating an illusion of movement. From this perspective, it is possible to analyze in a unified way the *Verlan* forms that appear to reverse syllables (mo.to → to.mo) and those which appear to reverse segments (fu → uf, luʃ → ʃəlu): in both types of cases, exactly one immediate precedence relation has been altered. Anti-IP is an anti-faithfulness constraint (Alderete 1999), and therefore co-exists with the corresponding ImmediatePrecedence (IP-Linearity). In *Verlan*, Anti-IP ranks over IP-Linearity, straightforwardly accounting for the verlanization of monosyllables (Tab.1).

Determining the size of the moved constituent The verlanization of disyllables appears to involve syllables permutation. However, the joint effect of Anti-IP, Max(C), Max(V), IP-Linearity, and the French phonotactics, does not rule out candidates in which an initial segment has moved by itself to the edge of the word (candidates c. and d., Tab.2). Illegal outputs are ruled out by Max($\widehat{\text{CV}}$) which requires that if a vowel has a transition with a consonant in the input word, it must have a transition with a consonant in the output. As a result, the apparent permutation of syllables is a consequence of a more general principle which maximally preserves $\widehat{\text{CV}}$ transitions. A welcome consequence of this analysis is that it accounts for disyllable verlanization without relying on the French syllabification, which wrongly predicts [mis.tik] ('*mystic*') → *[tik.mis] instead of [stik.mi]. Together with the assumption that medial clusters must be preserved in *Verlan* (CC-Contiguity), the present analysis derives the right output (Tab 3.).

(1) [fu]	Anti-IP	Max($\widehat{\text{CV}}$)	IP-Lin
> a. uf		*($\widehat{\text{Cu}}$)	*(f→u)
b. fu	*!		

(2) [bufõ]	Anti-IP	Max($\widehat{\text{CV}}$)	IP-Lin
> a. fõbu			*(u→f)
b. bufõ	*!		
c. õbuf		*!($\widehat{\text{Cõ}}$)	*(f→o)
d. ufõb		*!($\widehat{\text{Cu}}$)	*(b→u)

(3) [mistik]	Anti-IP	CC-Contig	IP-Lin
> a. stikmi			*(i→s)
b. mistik	*!		
c. tikmis		*! *(s→t)	

Predicting the rule-exception divide with lexically indexed constraints

Aleksei Nazarov, University of Massachusetts Amherst anazarov@linguist.umass.edu

Learning a grammar which finds a general rule while also accounting for exceptions is a long-acknowledged challenge (Chomsky & Halle 1968). I offer here an approach that replaces (stress) Faithfulness with lexically indexed constraints (Pater 2000, 2010). I show that this leads to an appropriate rule-exception balance without specifying a mechanism to achieve this, as opposed to existing Faithfulness-based learners, which need a Markedness-over-Faithfulness force (Smolensky 1996, Hayes 2004, Prince & Tesar 2004) or other mechanisms (Jarosz 2006).

My model is set up in Expectation Driven Learning (EDL, Jarosz submitted), which learns probability distributions over constraint rankings. Ranking probabilities are estimated by comparing, for each learning datum and constraint pair A and B, the relative accuracy of generating the datum for A >> B and B >>A given the current grammar.

I propose that a datum is declared exceptional for constraint B if that datum prefers B >> A, while previous data points in the data buffer on average prefer A >> B. Exceptions to constraint B trigger the induction of an indexed version (B_i) of that constraint. B_i and B may be ranked independently (e.g., B_i >> A >> B). All exceptions to A >> B are indexed to B_i unless there is evidence otherwise.

I test the learner on Dutch stress, standardly seen as quantity-sensitive (QS) with exceptions ((1a); van der Hulst 1984; quantity-sensitivity confirmed experimentally by Domahs et al. 2014). However, if all antepenult stress words are seen as exceptions, a quantity-insensitive (QI) rule covers the data as well ((1b); van Oostendorp 2012).

(1) a. *QS: antepenult if the word ends in Light-Heavy ((C)VC=Heavy); else, penult*
regular: [ˈva.ni.tɑs], [ka.ˈsi.no] ‘vanitas’, ‘casino’ exception: [ˈki.mo.no] ‘kimono’
b. *QI regular stress: penult; disconfirmed by Domahs et al. (2014)*
regular: [ka.ˈsi.no] exceptions: [ˈva.ni.tɑs], [ˈki.mo.no]

The learner saw a set of 36 constructed 3 and 4-syllable words with type frequencies of stress (antepenult, penult, final) by weight type (-L-L#, -L-H#, -H-L#) proportional to those in actual Dutch mono-morphemes (Ernestus & Neijt 2008). Constraints were taken from Nouveau (1994). Testing the learned grammars on previously unseen words, as in (2), reveals an unambiguously QS default pattern, albeit with some variation between 3 and 4-syllable words. Crucially, the QI hypothesis (van Oostendorp 2012) is rejected.

(2) *Number out of 10 runs in which the desired position for stress according to van der Hulst (1984) receives primary or secondary stress in novel words – per weight type and word length. Settings: 730 iterations, plasticity = 0.1, r = 50, data buffer = 25.*

	-Light-Light: penult	-Light-Heavy: antepenult	-Heavy-Light: penult
3 syllables	10	10	10
4 syllables	9 (1 antepenult)	10	3 (7 antepenult)

Thus, the proposed learner predicts quantity-sensitivity in Dutch by excluding certain antepenult stress words from exceptionhood, which is achieved through interaction between exception detection and EDL, rather than a dedicated mechanism.

THE ONTOLOGY OF ENGLISH MORPHO-PHONOLOGY

Heather Newell, UQAM
newell.heather@uqam.ca

This talk offers a novel analysis of the morpho-phonology of class 1 & 2 affixes in English. It is argued that two types of exceptional patterns normally attributed to class 1 affixes have two separate sources. Rendering these patterns non-exceptional requires that ***(i) English main stress is computed 'late' (after all cyclic phonology) and (ii) that 'class 1' affixes be recognized as infixes.***

The first exceptional pattern is due to the misclassification of a group of affixes, such as *-ology, -ee* and *-ese,* as class 1, due to their ability to affix to bound roots and to 'attract' stress. These affixes are, in fact, class 2 affixes. The crucial phonological data in of support this is: ***(1)*** *wom(ɦ)ology, gan(ɡ)ése, kin(ɡ)ée (in checkers)...* ***(2)*** *puzz[l̩]ése, twink[l̩]ology, butt[n̩]ée...* In (1) final consonants do not resyllabify as onsets to the V-initial affixes; we see NC cluster resolution. 'Normal' class 1 affixes, like *-ify* in *signify* or *-ard* in *bombard* bleed NC resolution by syllabifying final Cs as onsets, while class 2 affixes do not; *si(ɡ)ning, bom(ɦ)ish*. Root-final sonorants are also syllabic before misclassified *-ology, -ee*, and *-ese* in (2) despite the affixes being V-initial and therefore affording onset positions for resyllabification (c.f. *twink[l̩]ing* 'act of twinkling' vs. *twink.ling* 'a short moment' (Marvin 2001)). This leads to a paradox; *-ology*-type affixes are argued to be level 1 as they shift main stress (*bútton* vs. *buttonólogy*), but also must be level 2 to allow NC-resolution. If English main stress is post-cyclic, however, the paradox dissolves.

(3) **a. cycle 1:** [(bù.ttn̩)]$_{PWd}$ **b. cycle 2:** [[(bù.ttn̩)]$_{PWd}$ (ò.lo).gy]$_{PWd}$
c. main stress: [[(bù.ttn̩)] $_{PWd}$ (ò.lo).gy]$_{PWd}$ → [[(bù.ttn̩)] $_{PWd}$ (ó.lo).gy]$_{PWd}$

In (3a-b) footing is restricted to each cyclic domain. *'gy'* is extrametrical (following, ex. Hayes 1982), but the remainder of the class 2 affix, *'olo',* is large enough to form a moraic trochee (the long vowels, as in *-ee*, cannot be split, therefore extrametricality involving such affixes is illicit; they will also be footed). In (3c) the rightmost foot receives main stress. Note that level 2 affixes that are monosyllabic with short vowels (*-ment, less, ness, ful, ly*...), are phonologically small enough that their status as outside the stress domain derives not only from their affixation at cycle 2, but crucially from their inability to independently form a licit foot. Here, the rhymes of these affixes are extrametrical, and the affixes are sub-minimal. The prediction here is that there are no level 2 affixes that are monomoraic (after extrametricality) and attract stress, or that are bigger and do not attract stress. This is borne out by the data. **The second** exceptional pattern is the behaviour of true class 1 affixes when they attach, not to a root, but to another class 1 or 2 affix, the 'class 1 outside class 2' pattern being especially problematic (Halle & Vergnaud 1987). This is problematic within any framework that assumes a cyclic freezing effect (ex. Strict Cyclicity (Kaye 1995), PIC (Chomsky 2001)); an inner domain is altered by derivation in an outer domain. Crucially, this affixation order, while allowing the class 1 affix to affect the position of main stress, does not affect the foot pattern in cycle 1 (4a). The boundary between a level 2 affix and its base is maintained.

(4) a. cycle 1: [(prò.ba).ble]$_{PWd}$ **b. cycle 2:** [[(prò.ba).ble]$_{PWd}$ ist]$_{PWd}$
c. cycle 3: [[[(prò.ba).ble]$_{PWd}$ ist]$_{PWd}$ ic] $_{PWd}$ **d. infixation:** [[[(prò.ba).ble]$_{PWd}$(ìs.tic)]$_{PWd}$ic] $_{PWd}$
e. main stress: [[[(prò.ba).ble]$_{PWd}$ (ís.tic)]$_{PWd}$ ic] $_{PWd}$ f. **(prò).(bà.ble).(ís.ic)*

Upon affixation of class 1 *-ic* under a traditional analysis, the whole word is subject to restressing, which would result in an ungrammatical parse: (4f) (cf. *tìcònderóga)*. The correct pattern is predicted, and collapsed with the analysis in (3), under the proposal that class 1 affixes are infixes to the final segment of the morpheme to their left, and are merged into its phonological domain. It is this Phonological Merger (Newell & Piggott 2014) that qualifies as infixation. In (4c-d), *-ic* infixes into the prosodic domain of *-ist*. As *-ist* is outside the PWd containing *probable,* and is not an infix itself, re-footing of the inner string will not apply. Note that in this analysis both infixing (class 1) and non-infixing (class 2) affixes may root- or word-adjoin (Fabb 1988), so all root-attaching affixes will have a uniform (cycle 1) phonological behaviour, and non-isomorphism between phonological and semantic domains is correctly predicted (ex. *globe/global/globality*, where *globe* vs. *global* is semantically opaque, but *global* vs. *globality,* even though the entire word constitutes a single phonological domain, is not (cf. Lownstamm 2010). The proposal here is unique in accounting for the phonological, morphological, and semantic patterns evidenced by English morpho-phonology.

A Hierarchy of Phonological Emptiness

Marc Van Oostendorp - marc.vanoostendorp@gmail.com - Meertens Institute
Edoardo Cavirani - cavirani.edoardo@gmail.com - Meertens Institute

Main claim. We propose that not all empty nuclei (EN) are the same: some have more strength to license certain properties of neighbouring consonants than others. We submit that we can establish a hierarchy of phonologically empty positions which follows from a simple representational assumption. The more structure a position has, the stronger is its licensing power (if licensor) or the higher its need for licensing (if licensee). We demonstrate this on Dutch and Italian dialects. **Theoretical background**. If we assume that phonological analysis has the option of postulating EN, we have to observe that not all EN behave in exactly the same way. E.g. word-final EN allow for lenition of the preceding consonant in one dialect whereas in another they will not. Similarly, EN license complex syllabic structures in one dialect but only simple ones in another. Cyran (2008) and Scheer (2010) propose that we can describe this phenomenon by assuming that languages parametrically assign different degrees of licensing power. However, the taxonomy we can derive from these parameters is insufficient. E.g., schwa nuclei are not 'empty'; but at the same time they are arguably weaker power than other ('full') vowels. Also, the parameters are arbitrary: it is just assumed that empty nuclei have licensing power in some languages but not in others. Instead of this, we propose that licensing power corresponds in a 1-to-1 fashion to phonological structure. **Turbidity Theory (TT)** (Goldrick 2001, Revithiadou 2007, Van Oostendorp 2007) allows for autosegmental association to be split up into 2 asymmetrical relations: feature-to-segment (*projection*) and segment-to-feature (*pronunciation*). A feature can project to a segment in the underlying representation, but this does not guarantee that the feature is phonetically realized on this segment. In order for this to happen, phonology needs to match the projection with a pronunciation relation. Besides a way to deal with certain cases of opacity, TT also allows for a sensibly sophisticated view on a hierarchy of emptiness. The following segments are arranged in an order of ascending strength:

(1) a. x b. x c. x d. x
↑ ↑↓ ↑↓
@ @ A

In (1), a. is completely empty and therefore the weakest; b. has a feature projecting to it, so that it is stronger; c. entertains a fully symmetrical relation with the segment so that it is even stronger. However, since @ is itself a defective element, the relation is still not as strong as that with a 'full' A element (d.). **Processes**. Our proposal accounts for processes such as tone alternation and /d/ lenition in Limburg dialects (Van Oostendorp 2007), which is argued to display two 'phonetically' different word-final EN, the former being also phonologically empty (as in a.), the latter displaying instead an amount of phonological structure (as in b.) that is sufficient to trigger /d/ lenition (*róòd-róòj-róòjə* 'red.N/F/M'), but not to shelter a word-final obstruent from devoicing (*wíìs-wíìs-wíìzə* 'wise.N/F/M'). Other examples involve the empty 1S.PRES morpheme, which even blocks fricative devoicing and word-final /n/ deletion in St. Dutch, as well as /k/ epenthesis and /d/ lenition in Brussels Dutch. On the syllabic side, we show how a difference in the phonological make-up of 'phonetically' EN can account for the fact that Carrarese Italian licenses word-final T, R.T and TR structures (*cap* 'head', *colp* 'hit' and *libr* 'book') while the genetically and geographically related Pontremolese only T and R.T (*cap* 'head', *curp* 'hit' but *libar* 'book') (Cavirani 2015).

Testing OCP-like effects in infant word segmentation and learning

Mitsuhiko Ota and Barbora Skarabela
University of Edinburgh
mits@ling.ed.ac.uk; barbora@ling.ed.ac.uk

Recent experimental work has shown that adult speech processing is influenced by a constraint against within-word repetitions, which can be understood as a manifestation of the Obligatory Contour Principle (OCP; Leben, 1973; McCarthy, 1986). For example, in word segmentation tasks, Dutch speakers favor sequences without adjacent onset repetition (/potapa/) over sequences with such repetition (/popata/) (Boll-Avetisyan & Kager, 2014, 2016). Adult speakers may exhibit OCP effects because they have acquired the relevant phonotactic knowledge through generalizations over items in the native lexicon or from pre-lexical statistical patterns extracted from the phonological input. Alternatively, the OCP may reflect inherent constraints on speech processing that avoid repetitions in the signal. If OCP effects emerge from the statistical information in the phonological input or inherent processing constraints, we also expect infants to show similar effects. This study addressed this question by comparing infants' segmentation and learning of words that differed maximally with respect to phonological repetition.

In Experiment 1, we familiarized 9-month-old English-learning infants (N=24) to two passages, one containing a novel $C_1V_1C_1V_1$ word (e.g., /nini/, /fufu/) and one containing a novel $C_1V_1C_2V_2$ word (e.g., /nifu/, /funi/). During the test phase, we measured the infants' fixation times to a visual display, accompanied by auditory presentations of the familiarized words as well as a $C_1V_1C_1V_1$ word and a $C_1V_1C_2V_2$ word that had not been used during familiarization. Results showed a significant fixation time difference between the familiarized versus non-familiarized $C_1V_1C_1V_1$ words, but not between the familiarized versus non-familiarized $C_1V_1C_2V_2$ words, indicating that the infants were more likely to segment the reduplicated words out of the passages.

In Experiment 2, we familiarized 18-month-old English-learning infants (N=24) to two novel object-word pairs. One novel toy was labelled with a $C_1V_1C_1V_1$ word (e.g., /nini/, /fufu/) and another novel toy was labelled with a $C_1V_1C_2V_2$ word (e.g., /nifu/, /funi/). At test, we presented both toys and examined the direction of the infants' eye movement upon hearing the target word in a carrier phrase (e.g., "Where's the /nini/?"). Results revealed a faster and more reliable visual response to the toy with the $C_1V_1C_1V_1$ label than to the one with the $C_1V_1C_2V_2$ label. In word learning too, therefore, reduplicated structures were privileged over non-reduplicated structures.

To check whether this reduplication advantage may have roots in infants' exposure to phonological input containing a higher-than-expected probability of adjacent syllable repetition, we analysed the syllable transitional probabilities in the infant-directed speech in the Brent-Ratner corpus. The probability of immediate syllable repetition was not higher than the estimated chance level based on random combinations of all the attested syllables in the corpus.

Taken together, these results indicate that infants show a bias *for*, not *against*, parsing and learning words containing repetition of syllables, and that this bias is unlikely to be statistically extracted from the phonological input. Our results are most consistent with the interpretation that OCP effects emerge only through knowledge of the lexicon. Furthermore, the outcome of our study is in line with the proposal that humans are equipped with an experience-independent perception/memory bias that preferentially attends to repeated elements in the input (Endress, Nespor, & Mehler, 2009; Gervain & Werker, 2008). We argue that this bias does not figure in adult lexical representations because it is counteracted by a separate pressure on the lexicon, which militates against redundant phonological materials that do not contribute to differentiation of meanings.

Vowel reduction in Turkish

Markus A. Pöchtrager, Boğaziçi University, Istanbul

I. This presentation looks at vowel reduction in Turkish (TVR), with two central claims: (i) TVR is not a *phonological* process as per the criteria of Government Phonology (GP; Kaye, Lowenstamm & Vergnaud 1990, Kaye 1995) but (ii), far from being devoid of theoretical interest, TVR shows a complex interaction with and serves to signal morphosyntactic categories, similar to other phenomena in Turkish.

II. Alongside the standard/citation forms in (1a), colloquial Turkish has reduced forms (1b), cf. Lewis (1967), Vural (2006). TVR affects unstressed vowels preceding palatal *y* [j], *c* [dʒ]. Lewis claims (slightly incorrectly) that TVR only applies in *verbal bases*. (1) illustrates that it also affects *suffixes*; (2) the contrast to nouns.

(1)	a.	bekl**e**-yecek	anl**a**-yacak	kal-**a**cak	gel-**e**cek
	b.	bekl**i**-yecek	anl**ı**-yacak	kal-**ı**cak	gel-**i**cek
		‘wait 3SG.FUT’	‘understand 3SG.FUT’	‘stay 3SG.FUT’	‘come 3SG.FUT’

(2)	a.	boy**a**-yacak	~ boy**ı**-yacak	‘paint 3SG.FUT’
	b.	boy**a**-ya	~ *boy**ı**-ya	‘paint-DAT’

GP requires that for a phenomenon to count as phonological, it must be (i) exceptionless and (ii) non-arbitrary (i.e have a connection between target and trigger leading to loss/re-association of phonological material). Given the limitation to verbs (2), TVR cannot count as exceptionless. Non-arbitrariness is also problematic: The changes in (1) are expressible as the loss of the element **A**, taking *e*→*i* and *a*→*ı* (the realisation of an empty nucleus, cf. Charette & Göksel 1996), but it is unclear why that should happen next to palatal consonants characterised by the element **I:** There is no link between **A** and **I.** That result does not mean that TVR serves no function, cf. (3):

(3)	unreduced			reduced
	a. *Gel**e**ceğimi*	*söyledi*	*Gel**i**ceğimi söyledi*	
	come-FUT-1.POSS-ACC	say-PAST		‘S/he said I would come’
	b. *Gel**e**ceğimi*	*söyledi*	**Gel**i**ceğimi söyledi*	
	future-1.POSS-ACC	say- PAST		‘S/he told (me) my future’

Both times, *geleceğimi* is a form of the future participle of *gel-* ‘come’. The forms are case-marked, hence nominal. However, (3a) is the nominalisation of an entire phrase (VP/TP), but (3b) the nominalisation of the verb alone. Syntactic tests, e.g. whether modification is with adverbs or adjectives, support this. The contrast is replicable with any other verbal root, too. The (in)applicability of TVR functions as a signal.

III. TVR is not the only phenomenon of Turkish following a syntactic division. Firstly, the famous *k*-zero (=*ğ*) alternation (Inkelas 2011, Pöchtrager 2013) takes place in the majority of (but not all) nouns, but systematically fails to apply to verbs: *gerek* ‘necessity’/*gere**ğ**-i* ‘necessity-ACC’, but *gerek-* ‘to be necessary’/*gere**k**-ecek* ‘it will be necessary’. Secondly, Turkish has nouns with irregular stress (stress lexically marked non-finally), but no verbs with irregular stress (unless derived from irregularly stressed nouns). Thirdly, with the exception of one native root, there are no disharmonic verbs in Turkish (again, unless derived from disharmonic nouns), but plenty of disharmonic nouns.

GP often holds that (true) phonology serves parsing (Kaye 1989, Ploch 1996, Jensen 2000). Given the facts of Turkish, a wider perspective is called for. Patterns in phonological shape, though not true phonology, seem to signal morphosyntactic categories. In a language where the nouns/verbs/adjective distinction is sometimes claimed to be unspecified in the lexicon (Uygun 2009), this might be inevitable.

Prosodically-driven morpheme non-realization in the Minorcan Catalan DP

C. Pons-Moll (Universitat de Barcelona) & F. Torres-Tamarit (SFL, CNRS/Paris 8)

Minorcan Catalan has an intriguing case of morpheme non-realization, which is prosodically-driven and which cannot be accounted allomorphically. Kinship appositional phrases are generally realized with the structure Es-THE.M.DEF.ART conco-UNCLE en-THE.M.PERS.ART Jaume-JAMES ('uncle James') when the personal name starts with a consonant (1), but with the structure Es-THE.M.DEF.ART conco-UNCLE ∅ Àngel ('uncle Àngel') when the personal name starts with a vowel (2). That is, the personal article *en* fails to surface when the personal name starts with a vowel. The same path is detected for feminine personal names (3). (See (4), for the personal article paradigm in general conditions). Since the morphosyntactic structure in (1) and (2) is identical, the answer to this asymmetric behavior might be in the phonology. Our proposal is that the personal article has to be prosodically grouped with the preceding material (i.e. *es conco*), constituting a monolithic prosodic word preceding the one integrated by the personal name (as in ([əskòŋkun̪])$_{PWd}$ ([ʒáwmə])$_{PWd}$), and not with the following material (as in *([əskòŋku])$_{PWd}$ ([nán̪ʒəl])$_{PWd}$); a prosodic requirement which probably informs about a hierarchical structure among the constituents of the appositional phrase ([DP es [NP conco [DP en [NP Jaume]]]]) (*vs*. an alternative even-coordinated structure) and which might be equally adduced to explain the patterns of the feminine cases in (3). A morphoprosodic alignment constraint like ALIGN(R-PWd, Noun-L), which states that the left edge of the personal noun must be aligned with the right edge of the preceding prosodic word, outranking the constraint demanding the realization of a specific morpheme (*i.e.* REAL.-Morph.), explains the realization of the personal article in the cases of (1), and its lack of realization in the cases of (2). This constraint is vacuously satisfied in the cases of (4), where the personal name is not preceded by a prosodic word. Note how an external allomorphic account based on the double lexical representation /konku/ ~ /konkun/ (after a conceivable diachronic process of agglutination of *conco* + *en*) is not feasible: it is not possible to derive the selection of /konku/ before a word starting with a vowel, given the constraint ONSET (or, for the same purpose, *VV). One would expect, indeed, the selection of the alternative allomorph (*i.e.* /konkun/), which would entail the satisfaction of ONSET (or *VV) (cf. *es.con.**co. n'**Án.gel), but this is not the case (2). According to Selkirk (2001), there are two ways in which phonology may influence the morphosyntax of the sentence. Those where the phonological constraint ranking may force the non-realization (interpreted as deletion) of a word, and which are only possible if that deletion is recoverable (Pesetsky 1998). And those in which the phonological constraint ranking may force the non-realization of the whole sentence containing the function word, leading to a "crashing" of the derivation, and which are triggered when the deletion of the word is not recoverable. The cases exposed meet the first type of fate. When the adduced morphoprosodic conditions cannot be met, there is deletion of the personal article *en*. Deletion of *en* is possible because it carries redundant information and, thus, it is recoverable in Pesetsky's (1998) terms ("A syntactic unit with semantic content must be pronounced unless it has a sufficiently local antecedent"). The occurrence of another determiner in the DP and the fact that other varieties of Catalan do not use the personal article in the same environment (cf. *L'oncle Jaume*, *L'oncle Àngel*) seem to support this approach. In this talk, we are going to discuss these cases, and others where there is an actual crashing in the derivation (which leads to an alternative syntax), the interaction between the principle of recoverability and EVAL, and their consequences for the phonology-morphosyntax interface and the architecture of grammar. **(1)** es **conco en** [kòŋkun] Toni / Rafel / Pedro; l'**avi en** [àvin] Toni / Rafel / Pedro [avi = 'grandfather']; **(2)** es **conco** [kòŋku] ∅ Ignasi / Enric / Àngel; l'**avi** [àvi] ∅ Ignasi / Enric / Àngel; **(3)** sa **tia** ∅ Catalina / Margarita / Amparo 'F.DEF.ART. + AUNT + ∅ + PERS. NAME'; s'**àvia** ∅ Catalina / Margarita / Amparo 'F.DEF.ART. + GRANDMOTHER + ∅ + PERS. NAME'; **(4)** [ən] **en** Toni; [n] **n'**Ignasi; [nə] **na** Catalina; [n] **n'**Àngela 'PERS. ART. MASC. MASC. / FEM. FEM. + PERS. NAME'.

Converging evidence for biphonemic diphthongs in German

Renate Raffelsiefen & Fabian Brackhane, IDS Mannheim

Previous phonetic studies of the German diphthongs AI and EU, which contrast in words such as *Eile* 'hurry' vs. *Eule* 'owl', agree on recognizing two sequential target positions, connected by both raising and fronting. They further agree on the representation of the respective first targets ([a] for AI versus [ɔ] for EU), but not the second targets (e.g {[ɔœ],[ɔø],[ɔʏ],[ɔy],[ɔɪ], [ɔi],[ɔj]} for EU). Here we argue that the agreement is captured and the disagreement is resolved by recognizing the more abstract targets shown in (1). Significantly, these are precisely the segmental and prosodic representations motivated by strictly phonological evidence (see below). (The third diphthong, AU as in *Aule* ‚snot', is given in (1c).)

(1) a. 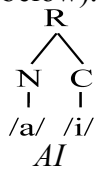

AI

b. 

EU

c. 

AU

The representations in (1) account for the properties mentioned above in that the consensus in identifying the first vowel is linked to the nucleus position, a position associated with maximal target approximation, whereas the disagreement is linked to the coda, a position associated with target undershoot (Krakow 1999). Drawing on these correlations between syllabic position and mode of phonetic implementation we conducted four phonetic studies to verify the role of /i/ as a target for the diphthongal offglides. We aimed to determine whether the phonetic realization of coda /i/ indicates a movement towards the phonetic target position observed for /i/ in nucleus position. To base the relevant projections for the offglides on solid ground we measured the formants F1, F2, and F3 in ten equidistant steps from 5% to 95% diphthong duration. Target positions for monophthongs like /i/ were measured in stressed syllabic position in words like *sieben.* The studies are briefly described below (I – IV).

I. Comparison of F1/F2-trajectories of AI and EU before various consonants (e.g. before /ç/ *Zeichen* - *Seuchen*), based on 50 male speakers of Northwestern Germany. Result: for AI-EU pair the projected F1/F2-trajectories consistently converge near the location of /i/. (Before -/f/ the projected trajectories meet in a point closer to /y/, due to regressive labialization.)

II. Comparison of F1/F2-trajectories of EU in the word *Zeug* in male and female speech, based on additional recordings of 50 female speakers of Northwestern Germany. Result: gender-specific differences between the trajectories observed in previous studies (Pätzold & Simpson 1997) follow directly from the different target positions associated with the monophtongs /ɔ/ and /i/ for females and males.

III. Comparison of F1/F2-trajectories of EU in the word *Leute* in read versus spontaneous speech based on 320 male speakers from all German speaking areas. Result: significantly closer proximity between the point at 95% diphthong duration for read (versus spontaneous) speech with the target position associated with syllabic /i/.

IV. Comparison of F3-F2 difference, known to indicate roundedness, for EU in the word *Leute* in read versus spontaneous speech based on 320 male speakers. Result: significantly stronger unrounding, in accordance with the unrounded target /i/, in read than in spontaneous speech. The last two results strongly support the presence of an actual target which is approximated the more closely, the more care and time is invested in the articulation.

The phonological motivation for the representations of the offglides with /i/ organized in the syllable margin concerns evidence (much of it novel) pertaining to markedness (e.g. the epenthesis of /i/ as a hiatus breaker in Old High German), neutralization (e.g. the absence of the palatal glide before AI and EU, but not AU, to satisfy a constraint against identical phonemes flanking the nucleus), and correspondence patterns (e.g. the prevalence of „impure rhymes" involving AI-EU in poetry, due to the same coda vowels in (1a,b), as opposed to the absence of AI-AU or AU-EU rhymes). It is the convergence of all types of evidence which supports a compelling argument for recognizing the biphonemic representations in (1).

Tobias Scheer CNRS 7320, Université Nice Sophia Antipolis

Intervocalic voicing is lenition (not spreading)

Problem: Lenition or assimilation, not both. Intervocalic voicing is known to be an instance of lenition: the intervocalic context is a typical lenition site. If so, intervocalic voicing couldn't be a case of assimilation, i.e. of the transmission of some phonological prime from the assimilating to the assimilated item. Lenition is positional and does not involve any transmission of primes: coda consonants for instance lenite no matter what the segmental environment, i.e. whether the preceding vowel is front, back, mid, high or low, and whatever the following consonant. Positionally defined lenition contrasts with assimilation, where instead some property is transmitted from a trigger to a target. It is therefore inconsistent to say that a process is an instance of lenition but in fact involves the transmission of some melodic prime (see Honeybone 2002: 205ff).

Existing analyses. This is especially true if, as is the case for vowels under typical assumptions regarding their internal structure, the allegedly donating segment does not possess the prime that it is supposed to transmit (here [voice]). In order to account for the participation of spontaneously voiced sonorants and vowels in the laryngeal phonology of non-spontaneously voiced obstruents, either an additional prime is introduced (Rice's 1993 [sonorant voice]) or featural representations are underspecified and [voice] is only introduced at a particular stage in the derivation (before or after obstruents have exchanged this prime, Itô & Mester 1986: 59f). These devices address the distinction between sonorant and obstruent voicing, but have nothing to say about the basic inconsistency mentioned: a process cannot be both lenition and assimilation.

Intervocalic voicing is (phonological) lenition plus (phonetic) contamination. The lenition cum transmission analysis is only inconsistent under the assumption (which is tacitly made by all approaches mentioned) that both processes are phonological in kind. Nothing withstands a scenario whereby phonology effects lenition and then the contamination of obstruents by vocalic properties occurs upon phonetic interpretation (spell-out from phonology to phonetics, cue constraints in Boersma & Hamann's BiPhon model). There is good evidence that intervocalic voicing is indeed lenition: for example, it typically occurs in a diachronic lenition trajectory that takes voiceless stops to fricatives or nothing in a number of steps, the first being voicing (e.g. Lass 1984: 178, Szigetvári 2008: 101ff). Since lenition is phonological in kind (nobody doubts that), the transmission of voicing in intervocalic voicing must not be phonological if the aforementioned inconsistency is to be avoided.

Analysis. Intervocalic voicing occurs in both voice and spread glottis languages. The analysis in the latter environment where voiceless obstruents bear [spread glottis] (or H) is straightforward: a phonological process of lenition eliminates laryngeal primes in intervocalic position, leaving a neutral C° ($p^{H}at^{H}a \rightarrow p^{H}at°a$) that goes to phonetic interpretation where (in spread glottis languages) C° receive a voiced interpretation anyway (t°↔[d]). In voice languages, the same phonological process eliminating laryngeal primes is active and now concerns intervocalic obstruents that bear [voice] (or L) ($p°ad^{L}a \rightarrow p°ad°a$). These delaryngealized items and lexically neutral C° are thus made phonologically identical: (p°at°a= p°ad°a): d° and t° are the same object (d/t only keeps trace of their lexical identity in the interest of exposition). Phonetic interpretation then is context-sensitive: neutral consonants C° are spelt out as voiced only in intervocalic position. That is, the neutral word-initial p° in p°at°a for example will continue to come out voiceless (unlike the intervocalic t° in this word and the delaryngealized d° in p°ad°a).

That delaryngealization is under truly phonological control and lenitive in kind is shown by the existence of patterns where obstruents in strong position are shielded against voicing. Western Romance is a case in point. First consider that in the evolution of this language family, like in many other cases, sonorants have the same influence as vowels and intervocalic voicing in fact is intersonorant voicing (where "sonorant" = vowels and sonorants). Hence voicing occurs in intervocalic contexts (V__V, Lat. RIPA > Fr. rive, Port. riba "river bank") and between a vowel and a sonorant (V__R, Lat. LIBRU, DUPLU > Fr. livre, double "book, double"). In strong post-consonantal position R__V, though, sonorants have no bearing on the voicing of obstruents: Lat. TALPA, VERSARE, CANTARE > Fr. taupe, verser, chanter "mole, to pour, to sing".

PROSODIC MISAPPLICATION IN COPY EPENTHESIS AND REDUPLICATION

Juliet Stanton & Sam Zukoff, MIT – {juliets,szukoff}@mit.edu

The term *copy epenthesis* refers to patterns of vowel epenthesis in which the featural value of the inserted vowel co-varies with context, "copying" the features of a neighboring vowel (e.g. /pra/ → [para], /pri/ → [piri]). This paper focuses on a class of cases in which the similarity between copy vowels and their hosts extends beyond featural resemblance, and how the existence of these effects informs the analysis of copy epenthesis – and by extension, the analysis of copying phenomena more generally. In particular, we show that copy vowels and their hosts strive for identity not only in all segmental features, but in all prosodic properties as well – and that this drive for prosodic identity can cause the misapplication of prosodic properties (i.e. stress, pitch, length). To explain these effects, we propose that copy vowels and their hosts stand in correspondence with each other (Kitto & de Lacy 1999). We show that this correspondence-based approach naturally extends to a class of similar misapplication effects in reduplication, and argue that the empirical overlap between the phenomena signals a formal similarity. In this way, the paper develops Kitto & de Lacy's (1999) suggestion that copying, phonological and morphological, is mediated by correspondence constraints (*cf*. Kawahara 2007).

COPY EPENTHESIS. One well-known example of prosodic misapplication in copy epenthesis comes from Selayarese (Mithun & Basri 1986, Broselow 1999, 2001). Stress is generally penultimate (*sampúlo* 'ten'), but epenthetic copy vowels (underlined) can cause stress to retract. When the copy vowel is word-final (*lámbere* 'long'), the antepenultimate vowel receives stress; when the copy vowel is word-internal (*karátu* 'card'), however, stress applies normally.

To explain these facts, we propose that copy vowels and their hosts stand in correspondence with each other (Kitto & de Lacy 1999). We can thus explain *lámbere*-type misapplication as the result of *stress matching*: copy vowels and their hosts must agree in stress (see also Kitto & de Lacy 1999 for a similar implementation). Assuming a prohibition on clashes (**lambéré*), the only way to ensure stress identity is to stress neither the copy vowel nor its host. In *karátu*-type forms, stress applies normally due to an inviolable ban on final stress (**karatú*): either remaining possible stress pattern will cause the copy vowel and its host to disagree in stress, so stress defaults to the penult. Below, we posit that the drive for copy and host vowels (subscripted with *x*) to agree in stress is enforced by IDENT[stress]-SC (SC = *"surface correspondents"*). Misapplication occurs in *lámbere* because IDENT[stress]-SC » *LAPSER ("*stress one of the final 2 syllables*"); it does not occur in *karátu* because NONFINALITY » IDENT[stress]-SC.

lamber	NF	ID-SC	*LAPSER
☞ a. lámbe$_x$re$_x$			*
b. lambé$_x$re$_x$		*!	
c. lambe$_x$ré$_x$	*!	*	

kartu	NF	ID-SC	*LAPSER
a. ká$_x$ra$_x$tu		*	*!
☞ b. ka$_x$rá$_x$tu		*	
c. ka$_x$ra$_x$tú	*!		

Beyond Selayarese, we show that positing a correspondence relation between copy vowels and their hosts easily accounts for similar stress-matching effects in Tahitian (Bickmore 1995), sub-categorical duration matching in Scottish Gaelic (Bosch & de Jong 1997), and a set of misapplication effects involving stress, duration, and nasality in Ho-Chunk (Winnebago; Miner 1979, *et seq*.). Kawahara (2007) argues against a correspondence-based analysis of copy epenthesis on the grounds that it predicts interactions between copy epenthesis and the assignment of prosodic properties. Since these effects are in fact attested, we argue that any adequate theory of copy epenthesis must posit a correspondence relationship between copies and their hosts. We discuss possible several reasons why this correspondence relationship might arise.

REDUPLICATION. Building on a suggestion by Kitto & de Lacy (1999), we explore possible formal connections between copy epenthesis and reduplication. In reduplication, like copy epenthesis, we find prosodic misapplication. For example, a number of languages display stress matching effects in reduplication, where base-reduplicant (BR) vowel pairs must match in the presence or absence of stress, and sometimes in stress degree (e.g. Ngan'gityemerri, Reid 2011; Diyari, Austin 1981; Indonesian, Cohn 1989). We argue that the existence of similar prosodic misapplication effects in copy epenthesis and reduplication points to an underlying structural similarity between the two phenomena: in both cases, correspondence constraints promote identity between members of a copying relation (Kitto & de Lacy 1999, *cf.* Kawahara 2007).

Another look at English phonotactics

Péter Szigetvári, Eötvös Loránd University, szigetvari@elte.hu

English (used here to mean standard British English) is traditionally analysed as having a complex vowel system comprising short vowels (eg *back*), long monophthongs (eg *bark*), and diphthongs (eg *bike*). Phonotactic constraints of English are formulated within this inventory: diphthongal offglides are not taken to be part of consonant clusters. The [j] of *yes* is analysed as a consonant, that of *say* as (part of) a vowel. My aim is to show that this prevalent view — the survival of the analysis of earlier stages of English — could be abandoned in favour of a much simpler system comprising no more than six vowels, [i], [e], [a], [o], [u], and [ə], potentially followed by [j], [w], or [ː], among other consonants (we here take the length diacritic to be a consonant).

Arguments against the complex-system view are numerous (detailed in Szigetvári 2015, in press): English is developing towards a system where any short vowel may be followed by any glide (eg *help* [hewp], *milk* [miwk]); we find epenthesis splitting the offglide of diphthongs and a following liquid (eg *fire* [faj⟨ə⟩(r)], *file* [faj⟨ə⟩l]); in some accents intervocalic lenition (=flapping) fails to occur after "diphthongs" and "long vowels" (cf de Lacy & Bye 2008), etc. In this paper I look at the phonotactic repercussions of analysing these complex vowels as vowel+glide sequences.

For data I use an online transcription dictionary, CUBE, which contains 100 000+ items with frequency counts taken from web searches, so not only type but also token frequency can be estimated in a rather large corpus: web contents indexed by Google in 2014. I have looked at the string of consonants after the last vowel in each item of the corpus, including forms containing the vowelless suffixes: past [d], pl/3sg [z] (and ordinal [θ], which happens to be missing in the most frequent clusters). The top dozen consonants and clusters ordered by token frequency are listed below.

_#	*toks*×10^9	*types*	_#	*toks*×10^9	*types*	_#	*toks*×10^9	*types*	_#	*toks*×10^9	*types*
j	81.12	10983	nt	11.24	1274	jtz	3.09	830	wsdz	0.18	26
n	49.21	7323	jt	9.77	1592	wsd	2.69	101	wntz	0.17	17
w	37.48	2061	jz	9.72	2934	ːsd	2.28	152	wndz	0.15	47
l	32.18	4350	nd	9.20	712	jlz	2.17	231	jntz	0.13	40
ŋ	26.60	6781	jn	8.45	1128	ntz	2.02	590	jldz	0.12	25
ː	25.82	1526	sd	8.10	1682	jmz	1.75	118	wnsd	0.10	11
z	24.67	5377	jl	7.77	548	ːtz	1.54	154	jndz	0.10	17
t	22.47	1778	jd	7.63	1035	jnz	1.41	372	ːsgd	0.08	6
d	20.78	3570	wz	7.26	832	jnd	1.33	209	ːnsd	0.07	9
s	18.76	3697	ːt	5.77	319	wld	1.28	141	jnʤd	0.05	12
m	13.62	1896	nz	5.65	2224	wnd	1.27	162	ːsgz	0.05	13
k	10.65	2728	kz	5.65	994	ŋkz	1.21	116	ːsdz	0.05	44

The data show that if the offglide of diphthongs (and long vowels) is taken to be a consonant, they do not clearly separate in their frequencies from other consonants. More importantly the maximal four-strong clusters include ones that do not contain a glide (though not in the top dozen shown here): [mptz] is #17 with 32.5 million tokens in 5 types, [kzdz] is #21 with 13.2 million tokens in 3 types, [ŋktz] is #24 with 0.8 million tokens in 3 types, etc. Ie the length of these clusters is constant only if we include diphthongal offglides in them. All of the four-strong clusters occur in morphologically complex words, shorter clusters may be morphologically complex or simplex independently of whether their first member is one of [j], [w], [ː], or some "real" consonant. These facts are interpreted as further evidence for the plausibility of taking all the diphthongs and long vowels of (standard British) English to be vowel+glide, thus dispensing with all complex vowels in the vowel system. So Trager & Bloch (1941) were right after all.

Structural Cumulativity in German Umlaut

Jochen Trommer, University of Leipzig – jtrommer@uni-leipzig.de

Summary: A central argument for the purely morphological status of German Umlaut (fronting of stem vowels in specific stem+affix combinations) is the fact that umlaut is not completely predictable, but exhibits significant subregularities (Köpcke 1988, Embick and Halle 2005), a pattern typical of 'soft' tendencies in the lexicon. In this talk, I show that the different arbitrary classes of roots and affixes postulated in morphological approaches such as Wurzel (1970) can be substantially simplified if they are reinterpreted purely phonologically as containing different amounts of floating vocalic features along the lines of Lieber (1987, 1992) and Wiese (1994, 1996). Combining different morphemes containing floating material leads to a gang effect for featural faithfulness constraints in Harmonic Grammar (Pater 2009). Umlaut thus provides a new type of evidence for cumulative effects involving only faithfulness constraints (Farris-Trimble 2008, Jesney 2015). **Data:** The basic empirical observation is that there are two classes of umlaut-inducing affixes, 'umlaut-enforcing' affixes as the diminutive suffix *-lein* which induces umlaut in virtually every stem and 'umlaut-triggering' affixes as adjectivizing *-lich* that do so only with specific stems, which in turn leads to a corresponding partition of stems in 'umlaut-prone' ones that exhibit umlaut before all umlaut-triggering affixes such as *Arzt* 'doctor' and 'umlaut-reluctant' stems such as *Amt* 'office' that only umlaut before umlaut-enforcing affixes.

	U-enforcing Affix	**U-triggering Affix**
U-prone Stem	Ärzt-lein	ärztlich
U-reluctant Stem	Ämt-lein	amt-lich

Analysis: Wurzel (1970) encodes these classes by diacritic morphological features, where appropriate combinations of feature values trigger word formation rules executing actual umlaut. I propose to capture Wurzel's original insights by encoding the propensity of morphemes to participate in umlaut by floating [-b(ack)] features on stems *and* affixes reconciling the positions of Lieber (1992) and Wiese (1996). U-prone stems and U-triggering affixes have one floating [-b] feature, umlaut-enforcing affixes two, and U-reluctant stems lack floating [-b]. Assigning to the constraint Ident [back] (protecting underlying backness of segments) a weight higher than that of Max [back] (which only applies to floating features, cf. Max Float in Wolf 2005, 2007), but lower than its multiples, predicts that a single floating [–back] is too weak to lead to realization (a), but any morpheme combination involving more than one instance of [-back] leads to overwriting since a single violation of Ident allows to avoid multiple Max violations which effectively cumulate (b,c,d) (affix material in blue).

a. **amt-lich Input:** = b.	Id[b] w=1.5	Mx [b] w=1	H
i. +b -b -b	-1		-1.5
☞ ii. +b -b -b		-1	-1

b. **ärzt-lich Input:** = b.	Id[b] w=1.5	Mx [b] w=1	H
☞ i. +b -b -b -b	-1		-1.5
ii. +b -b -b -b		-2	-2

c. **Ämt-lein Input:** = b.	Id[b] w=1.5	Mx [b] w=1	H
☞ i. +b -b -b -b	-1		-1.5
ii. +b -b -b -b		-2	-2

d. **Ärzt-lein Input:** = b.	Id[b] w=1.5	Mx [b] w=1	H
☞ i. +b -b -b -b -b	-1		-1.5
ii. +b -b -b -b -b		-3	-3

Finally I show that more complex differentiation between floating features and weights also allows to derive finer distinctions between morphological patterns capturing implicational relations

A stationary frequency effect in Manchester English

Danielle Turton[1], George Bailey[2], Maciej Baranowski[2] and Ricardo Bermúdez-Otero[2]
Newcastle University[1], University of Manchester[2]

The impact of lexical token frequency on phonetic implementation has been argued to support Exemplar Theory in the following way (Bybee 1998, 2002; Pierrehumbert 2001, 2002):

(a) Synchronically, high-frequency lexical items exhibit more coarticulation and reduction than low-frequency items (e.g. Dinkin 2008, Gahl 2008, Myers & Li 2009, among many others).
(b) This is because, in diachronic processes of lenition, frequent words change at a faster rate than infrequent ones.
(c) In turn, this is because high-frequency items suffer greater exposure to phonetic biases in production and perception than low-frequency items, and the effects of this difference are directly registered in phonetically detailed lexical representations.

This argument suffers from several problems. Hypothesis (b) has not been corroborated by actual diachronic observations in real or apparent time. Indeed, (a) does not logically entail (b): as acknowledged by Hay *et al.* (2015), frequent items can be ahead of infrequent ones, and yet change at the same rate. In such a scenario, the impact of frequency gives rise to a **constant rate effect** (CRE) in the sense of Kroch (1989): when modelled as logistic functions, the curves of change for high- and low-frequency items exhibit different intercepts but equal slopes. The existence of CREs in phonology was established by Fruehwald *et al.* (2013). Zellou & Tamminga (2014) report change in nasal coarticulation affecting high- and low-frequency items at the same rate. As regards (c), the empirical predictions of Exemplar Theory remain unclear. Sóskuthy (2014) shows that, in the absence of *ad hoc* stipulations, the inertia of a large exemplar cloud will cancel out the effects of greater exposure to phonetic bias. In addition, Hay *et al.* (2015) propose an exemplar-based account for a sound change apparently led by **low-frequency** words.

In this paper, we challenge (b) with evidence from a CRE in /t/-glottalling in Manchester. As expected, token frequency has a strong effect on /t/-glottalling, but there is no significant difference in the diachronic growth rates of glottalling in high- and low-frequency words. We demonstrate this statistically using LOESS-smoothers, mixed effects logistic regression, and Kauhanen & Walkden's (2015) mathematical model of the CRE. Our data come from a sociolinguistically stratified sample (62 speakers born between 1926-1985; 9,187 tokens of /t/ auditorily coded). Figure 1 (with data from word-medial /t/) shows that the curves of change in apparent time for high- and low-frequency items are not significantly different. Figure 2 shows the results of applying Kauhanen & Walkden's CRE model, which uses time-invariant contextual biases to derive context-specific curves from a single logistic growth function for all contexts (in this case, for all frequency bins). The model can be used diagnostically by comparing the error rates of CRE-constrained curves against independent logistic curves for each frequency bin. Fitting this more constrained model, with the CRE built in, leads to no increase in error over a model with completely independent logistic curves.

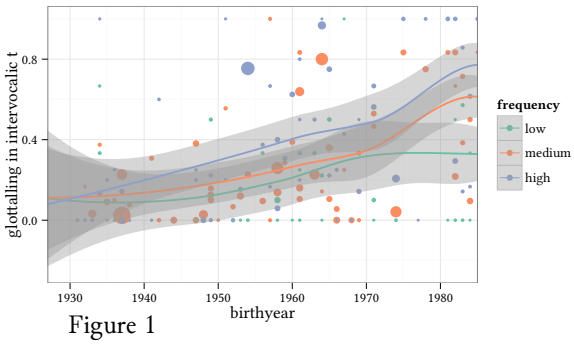

Figure 1

Further support comes from generalized mixed-effects logistic regression, which shows that an interaction between Zipf-scaled frequency (SUBTLEX-UK; van Heuven *et al.* 2014) and birthyear does not improve on a model without the interaction (by AIC or BIC). We conclude that the evidence stacks in favour of a scenario in which high- and low-frequency words change at the same rate, thus providing support for a CRE in Manchester /t/-glottalling.

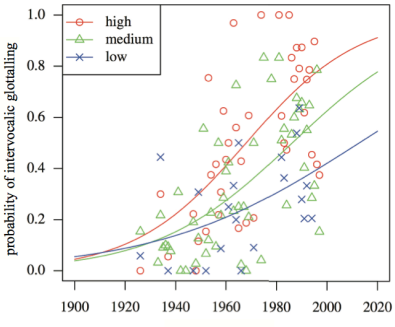

Figure 2

The absence of evidence for (b) suggests that alternatives to (c) should be considered. Frequency-driven CREs are consistent with modified versions of classical modular architectures in which neogrammarian innovation is effected through change in phonetic implementation rules referring to phonological categories in surface representations, whilst the impact of frequency is produced by orthogonal mechanisms (e.g. cascading activation, listener modelling).

Does "stabilization" entail de-opacification? The Korean evidence

Bert Vaux, Cambridge University (bv230@cam.ac.uk)
Kayeon Yoo, Cambridge University (ky259@cam.ac.uk)

Stratal OT generates opacity via inter-level constraint masking (Kiparsky 2000), predicting that if a process in stratum n that opacifies a process in earlier stratum n-1 moves into stratum n-1 itself (per the life cycle of phonological processes (B(ermúdez)-O(tero) 2007)), the interaction should become transparent. We argue on the basis of an experimental study of 12 speakers of Seoul Korean that although the process of Domain-Initial Denasalization (DID, Jones 1924, Cho and Keating 2001, etc.) has recently undergone "stabilization" (B-O and Trousdale 2012), moving into the phrase-level stratum, it still interacts opaquely with Regressive Obstruent Nasalization (RON, Kim-Renaud 1974 etc.), another phrase-level process.

RON spreads [+nasal] from a nasal consonant to an immediately preceding obstruent across Accent Phrase (AP) and weaker boundaries (Jun 1996; e.g. /guk-man/ [guŋman] 'soup-only'). This can be modeled in OT by ranking a *ON markedness constraint ("no obstruent-nasal sequences") between the IP and AP members of a family of IDENT constraints legislating nasality in domain-final obstruents: ID-NASO]$_{U}$, ID-NASO]$_{IP}$ >> *ON >> ID-NASAL O]$_{AP}$ >> … ID-NASAL. DID optionally (Kim 2011) targets nasal consonants in initial position of AP and higher prosodic constituents (Yoshida 2008), e.g. $_{U}$[$_{IP}$[$_{AP}${malhatɕa-mjʌn} $_{AP}${mʌri ap^{h}ɯta}]] → [balhatɕamjʌn bʌri ap^{h}ɯta] 'it gives me a headache when I talk about it'. This can be modeled in OT by ranking ID-NASAL between the AP and PrWd members of a parameterized markedness family punishing nasality in domain-initial position: *$_{U}$[N, *$_{IP}$[N, *$_{AP}$[N >> IDNAS >> *$_{PW}$[N, *$_{σ}$[N, *[N.

Before DID moved from the phonetics to the phrasal phonology, we'd have expected it to counterbleed RON in modular feed-forward models such as Stratal OT, (B-O and Trousdale 2012). Following DID stabilization, Stratal OT predicts that this interaction should become transparent (bleeding); e.g. for the input /…k} $_{AP}${m…/, the transparent [k b] is chosen over (and harmonically bounds) the opaque [ŋ b].

$_{AP}${hankuk} $_{AP}${mal...}	IDNASO]$_{IP}$	*ON	IDNASO]$_{AP}$	*$_{AP}$[N	IDNAS	*$_{PW}$[N
[…hankuk $_{AP}${mal…]		*!		*		*
[…hankuŋ $_{AP}${mal…]			*!	*	*	*
☞[…hankuŋ $_{AP}${bal…]			*!		**	
💣[…hankuk $_{AP}${bal..]					*	

Our experimental results reveal that in such cases speakers select the counterbleeding output [ŋ b]. This intra-level opacity cannot be attributed to "rule scattering" (B-O 2007) with both phrasal and phonetic versions of DID; this still predicts the incorrect bleeding interaction.

Although classic RBP (Kenstowicz 1994) can compute both transparent and opaque outcomes in such cases, it tends to opacity because rules are added to the end of their stratum (Chafe 1968, King 1973, Gress-Wright 2010). The results of DID stabilization in Korean suggest that our phonological architecture requires a mechanism such as RBP's (or OT-CC's) to produce intra-level opacity, consonant with synchronic arguments from Tiberian Hebrew (Green 2004), Sea Dayak nasal harmony, and Turkish k-deletion (Ettlinger 2009).

A New Typology of Tone and Stress Interactions

Draga Zec (Cornell University) and Elizabeth Zsiga (Georgetown University)

In this paper we investigate the phonological and phonetic structure of pitch accent in three Neo-štokavian dialects of Serbian (Belgrade (B), Novi Sad (NS) and Valjevo (V) idioms). These dialects are characterized by both lexical tone and stress (Lehiste & Ivić 1986, Godjevac 2000, Smiljanić 2002, Zsiga & Zec 2013). We argue for the phonological representation of pitch accent as a lexically associated High (H) tone, an association that is consistent across dialects. The distribution of stress is then predictable from the locus of H. Dialectal variation arises from the interaction of lexical tone, intonational tone and prominence, with different dialects giving precedence to different factors. We propose an Optimality Theory analysis, fully informed by both the phonetic and the phonological aspects of this prosodic system.

From a broader typological perspective, Serbian instantiates a hitherto undocumented type of tone/stress interactions. It is standardly assumed that stress and prominent (that is, H) tone preferably converge on the same syllable (de Lacy 2002, Hyman 2006). This according to de Lacy (2002) is captured by the foot as a representational device. The crucial ingredient is a set of OT constraints that enforce the association of H with the stressed syllable (head of the foot). Constraint *NONHD/H bans H from associating with the foot's non-head. Because this constraint is not matched with *HEAD/H, the result is an asymmetric typology that excludes languages with H on the non-head. However, it is precisely the configuration with H on the non-head that is found in Serbian. In (1) are shown pitch tracks for three trisyllabic forms, 'no$_H$ vine 'newspaper, 'mara$_H$ me 'scarves' and ra'mena$_H$ 'shoulders', with lexical H on the first, second, and third syllable respectively. Stress coincides with an initial H ("falling" accent), and falls one syllable ahead of a non-initial H ("rising" accent).

(1)

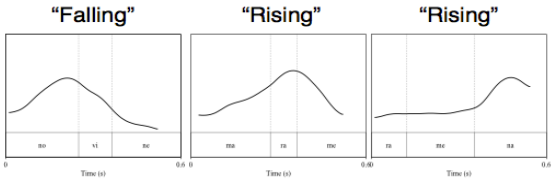

This pattern, shared by all three dialects, is analyzed as in (2, 3): input forms have only H tones, and stress in output forms is represented as a syllabic trochee. Initial H is on its head, as in (2a), and non-initial H is on its non-head, as in (2b,c). Thus H is preferably on a non-head but, with IDENTHIGH as an undominated constraint, never moves away from its source position. This pattern is captured by constraint ranking in (3), which operates at the lexical stratum. Presence of H within a foot is insured by HDPROSWD/H (Prosodic word's head foot is associated with H).

(2) a. no$_H$ vine → ('no$_H$ vi)ne b. mara$_H$me → ('mara$_H$)me c. ramena$_H$ → ra('mena$_H$)

(3) HDPROSWD/H, TROCHEE, IDENTHIGH >> *HEAD/H >> *NONHD/H

But while at the lexical stratum all three dialects prefer H on the non-head, they diverge postlexically. The results of an acoustic experiment show the following: In NS, pitch maxima correspond to the lexical locus of Hs in all positions; phrase-finally lexical H on the final syllable wins over intonational L%. In B, H is realized faithfully in all positions other than phrase-finally, where H on a final syllable retracts to the preceding *stressed* syllable under the pressure of intonational L%. In V, however, H retraction to the preceding *stressed* syllable is a more general process, which takes effect both phrase-initially and phrase-finally; H retraction is gradient, but is more intense phrase-finally, under the pressure of intonational L%. In sum, the realization of H is most faithful in NS, less so in B, and the least so in V (where gradience counterbalances this effect). NS and B differ in the ranking of REALIZE L%, with all lexical rankings remaining in place. In V, *NONHD/H is promoted postlexically above IDENTHIGH, resulting in H on the head.

Exceptional and derived environments in Assamese vowel harmony

Eva Zimmermann & Sören Worbs (Leipzig University)
eva.zimmermann@uni-leipzig.de, soeren.e.worbs@gmail.com

Main Claim Vowel harmony in Assamese involves exceptional triggers and derived undergoers. This complex pattern is taken to be a strong argument for morpheme-specific constraints in Mahanta (2012). In contrast, we argue that the Assamese exceptionality follows from general phonological repair processes in an account with floating features. That such a reanalysis in purely representational terms is possible hence weakens the need for morpheme-/construction-specific mechanisms in general (e.g. Pater, 2009; Inkelas, 2008).

Data Assamese (Mahanta, 2008, 2012) employs a pattern of regressive [ATR]-harmony (1-a) with [+ATR] as the dominant feature (1-b). The only low vowel [ɑ] is opaque and blocks the harmony (1-c). There are two exceptional contexts: *First*, the two adjectival suffixes /-ijɑ/ and /-uwɑ/ trigger regular [+ATR]-harmony but also an exceptional raising for an adjacent [ɑ] that then undergoes regular [+ATR]-harmony (2-a+b). *Second*, the exceptionally raised [ɑ] undergoes exceptional fronting if it is preceded by a mid front vowel (2-c). This fronting harmony applies only in these derived contexts of exceptionally raised vowels; underlying mid vowels are never fronted (/xɛh-ɔ-ti-ja/ → [xehotija], *[xehetija] 'recent', (Mahanta, 2012, 1112)).

(1) *ATR-harmony (Mahanta, 2012, 1112+19)*

a.	bɔx	-ɔ-ti	boxoti	'settlement'
	gʊl	-i	guli	'to mix'
	lɛk^h	-i-ka	lek^{h}ika	'writer'
b.	b^{h}ut	-ɛ	b^{h}utɛ	'ghost' (Erg)
c.	bɛpɑr	-i	bɛpɑri	'trader'
	zʊkɑr	-i	zʊkɑri	'shake' (Inf)

(2) *Exc. suffixes (Mahanta, 2012, 1121+32)*

a.	misɑ	-ijɑ	misolijɑ	'liar'
	bɔzɑr	-uwɑ	bozoruwɑ	'cheap'
b.	ɑlɑx'	-uwɑ	ɑloxuwɑ	'pampered'
	pɑtɔl	-ijɑ	pɑtolijɑ	'lightly'
c.	ɛlɑh	-uwɑ	elehuwɑ	'lazy'
	kɛsɑ	-uwɑ	keseluwɑ	'rawness'

Proposal We argue that the exceptional raising is the realization of a floating feature-geometric treelet (Bermúdez-Otero, 2012) containing [-low] and a V-place node that is part of the representation of the suffixes /-ijɑ/ and /-uwɑ/. The opacity of [ɑ] simply follows from undominated *+Low,+ATR excluding low ATR-vowels from the inventory of Assamese. In the context of the exceptional suffixes, however, the floating [-low] associates to the preceding vowel and hence makes it a licit target for the regular [ATR]-vowel harmony. The additional exceptional fronting applies since realization of the floating V-place node 'overwrites' the whole underlying feature specification of a vowel (3). It is shown how the ranking of standard faithfulness constraints predicts that this radically underspecified vowel then receives its [±back] feature from a preceding mid vowel. There is hence no additional process of exceptional progressive fronting harmony: spreading of [±back] is simply the optimal repair to provide a missing feature specification. That the exceptional raising is restricted to [ɑ]'s adjacent to /-ijɑ/ and /-uwɑ/ is a straightforward consequence from the standard assumptions that autosegmental association lines may never cross (Goldsmith, 1976).

(3) *Floating feature: Overwriting*

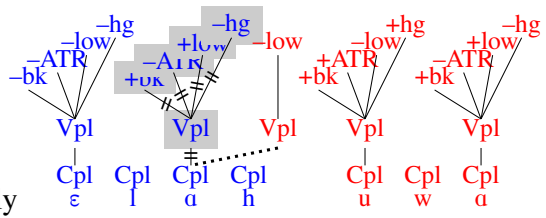

Discussion The alternative account for Assamese in Mahanta (2012) is based on morpheme-specific constraints that are restricted by an additional locality condition (Pater, 2009). We argue that this is not only dispreferred from the perspective of theoretical economy but also faces some concrete undergeneration problems for Assamese, especially concerning the additional fronting. The present analysis is based on rather standard assumptions about feature geometry and floating features (e.g. Lieber, 1992; Zoll, 1994; Wolf, 2007). That the latter not only account for mutation and/or ablaut but easily explain instances of apparent morpheme-specific phonology hence strengthens the empirical argument for floating features and weakens the arguments for morpheme-specific constraints and/or mechanisms.

Poster papers

Melodic primes as prosodic constituents

Phillip Backley and Kuniya Nasukawa
(Tohoku Gakuin University, Sendai, Japan)

This paper aims for a greater degree of melody-prosody integration by claiming that the basic units of segmental (melodic) structure also function as prosodic constituents, thus replacing traditional units such as onset, rhyme and foot. Using an Element Theory model of representation, it argues that elements (primarily |A|, |I|, |U|) not only express segmental contrasts but also combine asymmetrically to form complex structures which correspond to the domains of prosodic organization associated with constituents such as nucleus, rhyme, coda, syllable, foot and word. This denies these category labels any formal status, though the phonology can still refer to the relevant prosodic domains to describe phonological patterns.

Eliminating units such as 'onset', 'rhyme', 'coda' and 'syllable' from representations avoids the use of category labels that refer only to certain levels or domains of structure. After all, it is merely a stipulation to say that obstruents belong in onsets, or that sonorants belong in rhymes, since there is no formal link between obstruent/sonorant-type elements and the organizing nodes onset/rhyme. Instead, it is argued that the organizing categories traditionally labelled 'nucleus', 'rhyme' and 'syllable' are headed by a vowel element |A|, |I| or |U|, which may support dependent elements at various levels of embedding. The resulting concatenated structures – consisting only of elements related through head-dependency – are interpreted by language users as segmental strings of varying lengths from single 'syllables' to entire words.

The motivation for equating 'nucleus' with |A|, |I| or |U| comes from empty nuclei which, though melodically unspecified, may be realized as a language-specific vowel quality. This paper claims that a language's choice of default element/quality (|A|, |I| or |U|) is made by parameter, and that the chosen element functions as a structural head in an integrated melodic/prosodic structure: in English the default vowel is *ə* (i.e. the realization of bare |A| in a weak position), in Yoruba it is *i̵* (bare |I|), and in Japanese it is *ɯ* (bare |U|). These vowels create an acoustic baseline onto which other melodic properties are superimposed. Assuming the presence of a default head element in every language can explain why an 'empty' nucleus sometimes has the same phonetic quality as a lexically specified nucleus (e.g. Charette 1991; Harris 1994; Scheer 2004). In addition, it helps explain why default or epenthetic vowels are usually of a certain kind, namely [ə], [i̵] or [ɯ] rather than, say, [e], [æ] or [u].

The examples below illustrate the proposed approach. [i] has baseline resonance |A| (e.g. for English) with dependent |I|, which contributes its defining (marked) property.

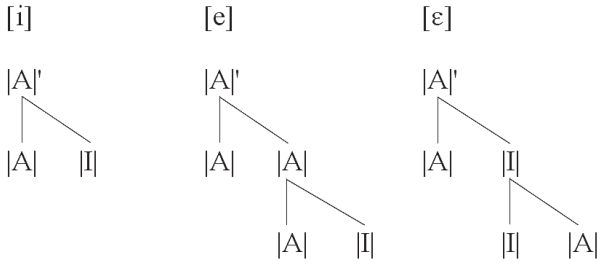

This way of concatenating elements asymmetrically is then applied iteratively, as in [e] and [ɛ], which differ only in their head category (|A| in [e], |I| in [ɛ]). Again, the most deeply embedded dependent element expresses the structure's distinguishing property (|I| in 'palatal' [e], |A| in 'open' [ɛ]). The phonetic value of an element compound is therefore determined not only by the elements it contains but also by the head-dependency relations between those elements.

A phonetic reflex of suffix in unsuffixed monosyllabic words with transparent vowels in Hungarian

Agnes Benkő, Eötvös Loránd Universty, Budapest, bagnes@infornax.hu
Ruben van de Vijver, Heinrich-Heine Universität, Düsseldorf, Ruben.Vijver@hhu.de

In recent studies it has been argued that small phonetic detail contributes to the expression of morphological categories (Gafos & Benus, 2006; Plag, Homann & Kunter, 2015). In this paper we investigate whether the frontness of a neutral vowel ([iː, ɪ, eː, ɛ] (Siptár & Törkenczy, 2000)) in monosyllabic Hungarian words correlates with their choice of suffix.

In monosyllabic words with transparent vowels the suffix class is unpredictable. Most monosyllabic words with transparent vowels take front suffixes, for example [viːz] *water* [viːz-nek] *water*-DAT, while some 60 words, such as [hiːd] *bridge* take back suffixes, [hiːd-nak] *bridge*-DAT.

A previous study has argued that this is the case. A transparent vowel in a monosyllabic word is fronter when the word takes front suffixes and more back when it takes back suffixes even when it occurs in isolation (Gafos & Benus, 2006). Gafos & Benus recorded the position of the tongue by means of EMMA and found that the tongue is more front in words that are inflected with front suffixes and more back when the word is inflected with back suffixes. Another study has argued that there is no such effect (Blaho & Szeredi, 2013). Blaho & Szeredi recorded monosyllabic words with a transparent vowel and measured the formants of the vowel. They compared words which take front suffixes with words that take back suffixes. On the basis of a mixed-effects regression analysis of the F2 they conclude that the transparent vowels are as front in words that take front suffixes as in words that take back suffixes, even though there is a tendency for transparent vowels to be fronter if they take front suffixes.

In order to shed further light on these contradicting interpretations, we conducted a production study in which we asked 21 participants to silently read a sentence with an inflected monosyllabic word with a transparent vowel. The participants had to read a second sentence in which the uninflected form of the test word was missing and had to pronounce the missing word. We then analyzed the formants of the transparent vowels.

Transparent vowels are indeed fronter in words that take a front suffix than in words that take a back suffix. In nonces (not shown) we found no such effect. Our findings support the view that there is a gradient phonetic influence on the realization of phonological –and in the end– morphological categories.

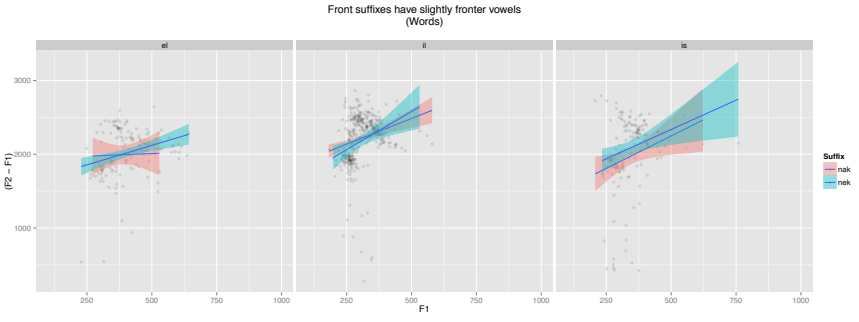

Figure 1: **The frontness (F2-F1) of transparent vowels in monosyllabic words. Transparent vowels in words that take front vowels are slightly fronter.**

Swabian pronominal variation

Tina Bögel
University of Konstanz ‖ Tina.Boegel@uni-konstanz.de

In contrast to standard German, Swabian, a major southern German dialect, allows for three different realisations of the first person nominative singular pronoun (1NomSg): [i:], [ə], and a pronoun drop. While the latter has been discussed in Haag-Merz (1996) and Bohnacker (2013), a complete postlexical phonological analysis of the pronoun drop and the distribution of all three forms has so far not been achieved. This paper proposes a solution with reference to information structure, the syntax–prosody interface, and postlexical phonology.

First, depending on focus, two of the variants can be distinguished: The full (in-focus) pronoun [i:] and the weak (unstressed) form [ə].

(1) jɛtst kɔx**=ə** ɛbəs vo: blo:s **i:** kɛn
Now cook.1SG.PRS=1SG.NOM something of.which just 1SG.NOM know.1SG.PRS
'Now I will cook something of which just I know.'

In (1), the unstressed [ə] is realised in the matrix clause. The fully stressed form [i:] is realised in the subordinate clause. While the syntactic positioning of the pronouns is not restricted in principle, there are restrictions on possible realisations based on prosodic constraints. The weak form [ə] is an enclitic and can thus not be realised without a host to its left. Consequently, constructions like in (2a), where the enclitic is in the initial position of the intonational phrase are not possible. However, if the enclitic follows a complementizer ((2b)) and is thus second in the corresponding intonational phrase, the construction is valid.

(2) a. *=ə kɛn di:
I know her
'I know her.'

b. das =ə di: kɛn
that I her know
'... that I know her'.

Second, as mentioned above, Swabian also allows for the pronoun to be optionally dropped (in contrast to standard German).

(3) vaɪʃ du: **vo:(=ə)=sə** nɔ̃: han
know.2SG.PRS 2SG.NOM where(=1SG.NOM)=3SG.F.ACC there have.1SG.PRS
'Do you know where I put her?'

The conditions for the pronoun drop are determined by postlexical phonology: The pronoun drop can only occur if a) the corresponding overt form is the enclitic [ə], b) the pronoun is part of a clitic cluster, and c) a valid syllable structure is preserved. For example, if pronoun drop would result in stray onset or coda elements, as shown in (4) for a stray coda, an invalid syllable structure obtains.

(4) *gɛʃtɐn hɐn**=Ø**=s ufgmaxt
Yesterday have.1SG.PRS(=1SG.NOM)=3SG.N.DAT open.PRF
'Yesterday, (I) opened it.'

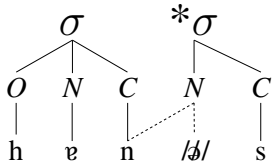

In terms of prosodic phrasing, the enclitic [ə] forms a nested prosodic word with its host. Evidence for this comes from an (optional) postlexical phonological process, *n-insertion*, which avoids a vowel hiatus between the host and the enclitic. The *n-insertion* neither occurs between two prosodic words, nor between two enclitics in a clitic cluster. It can, however, occur after a pronoun drop. As a result, it can be assumed that *n-insertion* applies after the pronoun drop and that it occurs at a one-level nested prosodic word: $((\text{host})_{\omega}\textbf{n-}{=}\sigma{=}\sigma)_{\omega}$. The clitic cluster in (3) thus has three possible variations: **vo:=n-ə=sə, vo:=ə=sə**, and **vo:=Ø=sə**, each derived from a set of ordered postlexical phonological processes: prosodic phrasing, syllabification, pronoun drop, and *n-insertion*.

Redistribution of stress-related prominence in a Belarusian dialect

Lena Borise | borise@fas.harvard.edu | **Harvard University**

Introduction. The main physical correlates of stress in Belarusian are greater duration, intensity and pitch of the stressed vowel as compared to the neighbouring unstressed vowels (e.g., Sussex & Cubberly, 2006:179). The stress system of the dialect spoken in the villages of Malyja Aŭciuki and Vialikija Aŭciuki in south-east Belarus has been reported (Kryvitski 1959, Belaja 1974) to exhibit a typologically unusual property: in the environments in which the pretonic vowel (V_1) is mid-low or low (/ɛ, ɔ, a/), and the stressed vowel (V_2) is mid-high or high (/i, ɨ, u, e, o/), V_1 exhibits equal or greater duration, intensity and pitch, as compared to V_2 (1a). On the other hand, in cases where V_2 is lower than V_1 or of equal height, V_1 exhibits significantly lower duration, intensity and pitch, as expected in Belarusian (1b). I will be referring to this phenomenon as **pretonic prominence (PP)**.

(1) a. *sestru* 'sister.ACC' [sʲ**ɛ:**ˈstru] b. *sestra* 'sister.NOM' [sʲɛˈstra]

Similar developments have been reported for some Russian dialects, such as certain Mosalsk (Broch, 1916), Vladimir-Volga Basin (Avanesov, 1927) and Tver (Nikolaev, 2009) dialects. The Aŭciuki phenomenon, however, differs from the above in that the changes that affect V_1 are conditioned by the height of V_1 and V_2 instead of applying across the board.

Since greater vowel duration, intensity and pitch are thought to correlate with stress, the question arises about the location and nature of stress in the Aŭciuki dialects. This paper reports on acoustic data obtained in Aŭciuki in 2015 by the author, recorded with a Zoom H4n recorder, and puts forward a preliminary analysis of PP.

Data. Data from three female informants (natives of Aŭciuki, > 60 yrs, M=72 yrs) was used. Seventy five tokens containing conditions in which PP is predicted to apply (a low or mid-low V_1 followed by a high V_2) were compared with twenty five tokens with no conditions for PP (both vowels low). All examples were extracted from the middle field of declarative clauses with all-new intonation. Highest values for intensity, pitch, and duration were extracted for vowels in four conditioning environments: (1) V_1, unmarked context; (2) V_2, unmarked context; (3) V_1, PP context; and (4) V_2, PP context. The values for duration, pitch and intensity of V_1 in PP contexts were compared with those of (i) V_2 in PP contexts, and (ii) V_1 in unmarked contexts.

Results. In PP contexts, V_1 is significantly longer (m_{V1}=114.9 ms, m_{V2}=86.8 ms; Wilcoxon paired test, $p < 0.01$) and higher in intensity (m_{V1}=76.3 dB, m_{V2}=72.5 dB; Wilcoxon paired test, $p < 0.01$) than V_2, while the pitch values of the two vowels are comparable (m_{V1}=231.8 Hz, m_{V2}=230.0 Hz; Wilcoxon paired test, $p > 0.05$). This picture is quite different from the distribution of acoustic prominence between V_1 and V_2 in the unmarked cases, where V_2 is significantly longer than V_1 (m_{V1}=75.2 ms, m_{V2}=135.0 ms; Wilcoxon paired test, $p < 0.01$), but both vowels are close in pitch (m_{V1}=198.0 Hz, m_{V2}=203.0 Hz; Wilcoxon paired test, $p > 0.01$) and intensity (m_{V1}=74.0 dB, m_{V2}=74.4 dB; Wilcoxon paired test, $p > 0.05$). Finally, all three characteristics of V_1 in unmarked contexts differ significantly (Wilcoxon test, $p \leq 0.01$) from those of V_1 in PP contexts.

Analysis. The domain of stress realization in the Aŭciuki dialect is disyllabic, encompassing both V_1 and V_2 -- cf. Williams, 1999 for a similar account of Welsh stress. That is not to say that both syllables bear stress, or that it can fall on either of the vowels within the domain. While stress stays in its etymological position, V_2, its acoustic correlates prefer to align with the more sonorous vowel of the two - the lower one. Specifically, in those instances where the two vowels, V_1 and V_2, are unequal in height, the lower one of the two will attract the acoustic prominence associated with stress - that is, higher intensity and longer duration, as well as high pitch. This account correctly predicts that in cases where V_2 is more sonorous than V_1, V_1 will not be affected.

Local vs global evaluation of foot-driven tone processes

Jeroen Breteler · j.m.w.breteler@uva.nl · University of Amsterdam

This presentation focuses on bounded and unbounded tone shift and spread, mainly reported on in Bantuist literature. Example (1) shows data from Rimi, where surface tone, marked with an acute accent, appears one syllable to the right of its sponsor location, which is underlined (Schadeberg 1979; Myers 1997). This is a case of bounded, rightward tone shift. Example (2) shows unbounded, antepenult spreading in Xhosa (Downing 1990).

(1)
a. ra-mú-ntu 'of a person'
b. mu-tem-í 'chief'

(2)
c. ú-kú-shúkúmisa 'to shake'
d. ú-kú-námáthélisa 'to cement'

Previous work has established that foot structure offers a unified account for these patterns, acting as a bounding domain in bounded patterns, and as a target for association for long-distance tone movement in unbounded patterns. However, investigation of the factorial typology revealed dubious predictions, specifically two types of surface tone gaps, and rampant tone-to-edge-proximity effects. A likely cause for this overgeneration are the tone-foot licensing constraints, exemplified in (3) (Zoll 1996; De Lacy 2002):

(3)

License(H, Foot-Right)	Assign one violation mark for each H tone that is not associated to the rightmost syllable of some foot

These constraints do double duty, acting both as the general drive for tone association, and as the targeting mechanism for attracting tone, causing it to spread.

In this talk I consider an alternative account, where Tone Attraction is evaluated not locally in the licensing context, but Globally across the entire prosodic domain ('GTA'). Finding the best formalization for such a tone attraction constraint is non-trivial, and is one of the subjects of the presentation. A tentative definition is shown in (4) (Martínez-Paricio and Kager 2015):

(4)

GTA-Right	Assign one violation mark for each toneless syllable that is not in an unbroken sequence of toneless syllables starting from the left edge of the word

This constraint can gradiently pull tones to the right, thereby obviating the need for gapped tones. The typological predictions of GTA were tested for a small set of inputs using OTWorkplace (Prince et al. 2015). In addition to accounting for the relevant attested patterns, GTA cuts down on edge effect patterns, shrinking the size of the results set by roughly 50%.

Interestingly, GTA may be unable to deal with more complex attested patterns; I will argue that GTA has trouble accounting for the full facts of Copperbelt Bemba, a language that combines bounded and unbounded spreading (Bickmore and Kula 2013).

In summary, GTA improves on previous typological work, accounting for attested patterns while avoiding some types of overgeneration and the use of tone gaps. Furthermore, GTA raises questions about what is allowed in constraint formulations, and reveals which languages stand out from the rest in terms of complexity (cf. Jardine 2015).

L1 Tone attrition among bilinguals in an L2 speaking environment

Xiangjie Cao xiangjie.cao@ncl.ac.uk

Newcastle University

Late bilinguals who continue to use their native language while using an L2 every day and/or residing in the L2 community have been shown to exhibit changes in their L1. The majority of the research on changes in L1 use and possible L1 attrition has focussed on the lexicon, morphology and syntax (Schmid 2002), but in recent years, attention has moved to phonology. Tonal attrition has received the least attention.

In Mandarin, tone is used to differentiate lexical items or to express morphological functions. There are four tones in Mandarin: the level first tone (T1), the rising second tone (T2), the falling-rising third tone (T3), and the falling fourth tone (T4). Among these, the tone considered to be the most complex is T3. Tone sandhi also applies to T3 where for two adjacent T3s, the first T3 is realized as T2 (Yip 1980). T2 and T4 show tone variations with different tones followed. In trisyllabic sequences, the middle T2 changes to T1 if the first syllable is T1 or T2 and the final syllable is a random tone from four tones.

Several studies over the past decade of Mandarin bilinguals have revealed attrition of tone by L1 Hakka Chinese speakers living in a Mandarin-speaking area (Yeh, 2011). Little is known, however, about what happens when a tone language speaker moves to a non-tone language environment.

The present study addresses whether there are changes in tone production and perception by Mandarin speakers living in a non-tone language speaking environment (the UK) for varying lengths of time. The study compares 50 Mandarin-English late bilinguals who had been living in the UK from three months to more than five years with Mandarin monolinguals (only with minimal English exposure at school) living in mainland China. Their perception and production of four tones at word and sentence level were tested by a listening comprehension task, an interview task, and a story-telling task for both formal and more casual speech. A questionnaire collected data on speakers' use of and contact with both languages.

The data were analysed acoustically using Praat (version 5.4.22) speech analysis software (Boersma & Weenink 2015), and statistical measurement revealed that late bilinguals who had lived in the L2 environment for over five years showed signs of attrition on T3, tending to omit the raising part in production, and the first T3 in tone sandhi. The bilinguals' four tones showed a tendency to merge, rendering them less distinctive than the control group's tone production. Age of arrival, amount and type of L2 exposure and of L1 contact showed correlations with tone attrition. Moreover, some patterns mimic tone acquisition (Li and Thompson 1976; Lin 1985; Chang 2014) indicating that markedness plays a role in both acquisition and attrition.

Logical characterizations of local vs. long-distance phonology

Jane Chandlee (jchandlee@haverford.edu), Haverford College, Haverford, Pennsylvania, U.S.A.

This paper presents a method for representing phonological maps with logical formulae over graph transductions and argues that a benefit of this approach is that both local and long-distance dependencies can be described without increasing the computational complexity of the formalism. The crucial difference is only whether the formula references an immediate predecessor or a general precedence relation among segments in the string.

As an example, consider nasal agreement, which is attested both locally, as in (1a) (Ganda, Ashton 1954, Choti 2013), and long-distance, as in (1b) (Kikongo, Rose & Walker 2004).

(1) a. **nl**imi ↦ **nn**imi 'tongues' b. tu**n**iki**d**i ↦ tu**n**iki**n**i 'we ground'

The goal is to define logical formulae that establish the conditions under which a segment surfaces as nasal. This is done by first representing the input and output strings as graphs, as in (2). These graph representations allow us to refer to particular positions in the string as well as segment correspondence between the input and output.

(2)

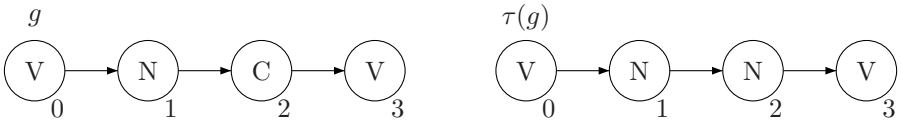

The graph g pictured on the left in (2) (the 'input' graph) is defined with 4 vertices, numbered 0-3, with labels V, N, C, and V, respectively, and 3 edges. The graph on the right in (2) (the 'output' graph) is the image of g under the phonological map of local nasal harmony shown in (1a). This graph can be defined in terms of g using the following formulae (following Engelfriet and Hoogeboom 2001).

(3) a. $NC(x) = (\exists y)[N(y) \wedge C(x) \wedge y \lhd x]$ b. $\varphi^0_N(x) = N(x) \vee NC(x)$
c. $\varphi^0_C(x) = C(x) \wedge \neg NC(x)$ d. $\varphi^0_V(x) = true$

Informally, (3a) evaluates to true iff there exist two positions in the input string, one labeled N and one labeled C, such that N immediately precedes C. (3b) asserts that any position x in the output string is labeled N iff its input correspondent is labeled N or $NC(x)$ is true. (3c) asserts that any position x in the output string is labeled C iff its input correspondent is labeled C and $NC(x)$ is false. Lastly, (3d) requires positions labeled V in the input to also be labeled V in the output.

The advantages of this approach include the following. First, long-distance maps like (1b) satisfy the exact same formulae in (3), provided the $y \lhd x$ relation is defined as general *precedence* rather than immediate predecessor (Heinz 2009, 2010). Thus both local and long-distance phenomena are accounted for by changing the relational structure rather than increasing the power of the formalism. Second, beyond string representations, logical formulae can apply to other phonological structures, such as autosegmental representations (Jardine 2016) and metrical trees. As above, describing phenomena in these domains is a matter of changing the representation, not the power of the logic. And lastly, as the computational complexity of various types of logic is well-studied (Rogers & Pullum 2011, Rogers et al. 2013), identifying the most restrictive logic that is descriptively adequate for phonological maps will lead to a better characterization of the range of 'possible' maps.

Word- vs. stem-level effects in French vowel alternations

Adeline Charlton • Newcastle University • a.charlton5@ncl.ac.uk

The French sandhi phenomena of liaisons have long been debated in phonological research (see Durand, 2008 or Côté, 2010 for a review), yet few studies, and only recently, have focused on variation patterns of the vowel preceding the liaison consonant (Eychenne 2014, Bermudez-Otero, 2015) highlighting the mixed nature of these vowels with features from both forms (masculine/feminine or citation/concatenated).

This paper aims to contribute additional data on vowel alternations to inform further accounts of liaisons phenomena, and to provide empirical evidence in support of stratal OT. Consider intervocalic alternations of the nasal vowels of French:

(1) Ṽ ~ ṼNV ~ VNV

Prosodically, these patterns match phonological domains only if one ascribes to a richer prosodic hierarchy, i.e. admitting an intermediate level between phrase and word. Although, this richer hierarchy is still debated today, further evidence (Downing, 2015) shows that some phenomena are limited to certain domains and therefore cannot be captured by an obscuring or overlapping label such as PWord, thus defending the idea of an extra constituent, namely PStem.

Focusing more specifically on French nasal vowel intervocalic alternations reveals that changes in vowel features and stress patterns motivate a division word vs. stem, where the phonology of each level accounts for the distribution pattern in (1):

(2)

Level				
Phrase	ṼV	ɑ̃ˈplɛ̃alˈbi	*en p**lein A**lbi*	‘right in the middle of Albi town’
Word	ṼNV	ɑ̃ˈplɛ̃neˈte	*en p**lein été***	‘right in the middle of the summer’
Stem	$V_{[+lo]}$NV	ɑ̃ˌplɛneˈte	*en p**lein été***	‘right in the middle of the summer’
	$V_{[-lo]}$NV	ɑ̃ˌpleˈnɛʁ	*en p**lein a**ir*	‘outdoors’

A stratal account elegantly captures the role played by cohesion and semantics, often invoked in an ad-hoc fashion to define liaisons, by simply embedding these notions in its stratified phonology principle. Moreover, by allowing recursivity at the stem level, it also accounts for variation in closed syllable adjustment (CSA) or *loi de position*, for example between in ɛ ~ e in (2).

In doing so, a stratal account also incorporates previous analyses of clitics (among others Anderson, 2005). That is, assuming that the prenominal adjectives in (2) are clitics, as affixes to the word level they show the least incorporation into the word (free clitics), and the most complex feature combination, whereas as affixes to the stem level, they display a variable integration into the stem (affix or internal clitic) with a simpler feature combination. In this approach, alternations of nasal vowels can be seen as varying degrees of lexicalisation or, in stratal terms, various degrees of stratification, yielding variable feature combinations.

Cortical dynamics explain language grammar: Insights from a latching Potts network

Joe Collins[1]

1. CASTL, University of Tromsø (Norway), joe.collins@uit.no

This study employs a complex dynamical systems approach (Potts-model) to investigate aspects of natural language processing.

Linguists have long observed that natural languages employ a variety of strategies to avoid repetition of similar segments (e.g. English *bus* [bʌs] + plural – *s* [-z] gives *buses* [bʌsəz]). Why this should be a feature of natural language has largely remained a mystery.

Previous studies using Potts-models have shown that a combination of neuronal fatigue and rapid inhibitory mechanisms will cause the network to "latch" between discrete memories stored in the system (Treves 2005). The strings resulting from latching have been shown to exhibit grammar-like properties, dependent on the overlap in the representations of the memories (Pirmoradian and Treves 2012). This study argues that the same mechanisms can account for the *repetition-avoidance* of natural language.

A latching Potts network was created which represents speech segments as distributed patterns of activity. The refractory period created by rapid inhibition prevents a currently active population of neurons from remaining active during the subsequent state. Therefore, if similar speech segments have strongly overlapping representations, the network can never latch from one to the other.

The representational assumptions are congruent with current electrophysiological data on the encoding of speech information in the Superior Temporal Gyrus (Mesgarani et al. 2014). The network model is therefore hypothesized to simulate an extended network of long-range connections between the STG and premotor areas.

Mesgarani, N., Cheung, C., Johnson, K., Chang, E. F. (2014) Phonetic feature encoding in human superior temporal gyrus. *Science 28*: 343, 1006-1010.

Pirmoradian,S. Treves, A. (2012) A talkative Potts attractor neural network welcomes BLISS words. *BMC Neuroscience 13*:P21.

Treves, A. (2005) Frontal latching networks: a possible basis for infinite recursion, *Cognitive Neuropsychology 22*:3-4, 276-291.

The impact of frequency on pattern application

Bartłomiej Czaplicki
University of Warsaw, bczaplicki@uw.edu.pl

Three word formation processes from Polish illustrate the impact of frequency. In the first one, type frequency determines which of the two competing patterns is productive. In the second process, the applicability of base transparency is mitigated by token frequency. While transparent bases are preferred in rare and novel words, mutations are permitted in well-entrenched words. The third process, showing the highest type frequency of the three, invariably results in mutations.

Two patterns (schemas) of noun formation with *-arz* are available in the same phonological context. One shows consonant mutations, the other does not.

(1) a. dru[k] ~ dru[k]-arz
mle[k]-o ~ mle[ɡ̑]-arz

b. mły[n] ~ mły[n]-arz
komi[n] ~ komi[ɲ]-arz

The reliability of a morphophonological schema is a function of its frequency and is defined as a ratio of HITS (the number of words that it applies to) to SCOPE (the number of words that meet its structural description) (Albright & Hayes 2002).

(2) a. […k] ↔ [[…k]aʒ] 80/86 ⟹ 0.93
[…k] ↔ [[…ɡ̑]aʒ] 6/86 ⟹ 0.07

b. […n] ↔ [[…n]aʒ] 5/53 ⟹ 0.09
[…n] ↔ [[…ɲ]aʒ] 48/53 ⟹ 0.91

Novel words show the application of only one of the available schemas, the one with the highest reliability score, e.g. *kaja*[k] ~ *kaja*[k]-*arz* and *ochro*[n]-*a* ~ *ochro*[ɲ]-*arz*.

Nouns in *-ist-a/-yst-a* fall into three categories with respect to the presence or absence of mutations between the base and the derivative.

(3)

a.	non-alternating pattern	Bonapar[t]-e	bonapar[t]-yst-a
b.	vacillating words:		
	non-alternating pattern	al[t]	alt[t]-yst-a
	alternating pattern	al[t]	alt[ʦ]-ist-a
c.	alternating pattern	fle[t]	fle[ʦ]-ist-a

It is argued that these three different patterns of behavior can be explained by reference to token frequency and its interaction with two factors: the robustness of morphophonological schemas (gauged by type frequency) and the preference for transparent bases. The presence/absence of mutations in the base reflects the degree of entrenchment of a particular word, with rare words showing transparent bases and established words respecting regular morphophonological schemas. Words of medium frequency exhibit fluctuations. Hay's (2003) dual-route model of lexical access provides useful tools in explaining this patterning.

In contrast, diminutive formation using the suffix *-ek* always leads to mutations for base-final velars, e.g. *kro*[k] – *kro*[ɡ̑]-*ek*, **kro*[k]-*ek*. The impact of base transparency diminishes as the robustness of a morphophonological schema increases. This agrees with the facts, as fluctuations are not attested for robust patterns.

The three processes are given an analysis in terms of OT constraints, where schemas, ranked along a continuous scale according to their frequency, interact with other pressures derived from language use: the strength of mental representations (gauged by token frequency) and base transparency.

Codas and branching onsets during French phonological acquisition

Laetitia de Almeida, Sandrine Ferré & Christophe dos Santos
Université François-Rabelais, Tours, & Inserm U930
laetitia.dealmeida@univ-tours.fr; sandrine.ferre@univ-tours.fr; christophe.dossantos@univ-tours.fr

The studies conducted in phonological acquisition generally agree that branching onsets and codas are among the last syllable constituents acquired by children. This has been noticed in typical first language (L1) acquisition (*e.g.* Fikkert 2007), as well as in bilingual acquisition, whether simultaneous (2L1) (*e.g.* Kehoe & Lleó 2003) or sequential (L2) (Rattanasone & Demuth 2014). On the other hand, branching onsets and codas also constitute a challenge to children with specific language impairment (SLI), whether bilingual or not, who have trouble producing these structures correctly at a later stage of development (Ferré et al. in press). Typically, most authors point to branching onsets being even more problematic for children than codas. Nevertheless, most of the studies do not make a distinction between internal and final codas, assuming that all final consonants are codas. However, following Piggott (1999)'s proposal, word-final consonants can be analyzed as codas or as Onsets of Empty-Headed Syllables (OEHS) depending on the phonological properties of the target language. This has clear implications for data analysis, especially in languages in which word-final consonants are claimed to be OEHS, like in French (Kaye, Lowenstamm & Vergnaud 1990): in order to compare the acquisition of branching onsets and codas, one should focus on internal codas only. In fact, some studies conducted on French phonological acquisition, whether in a monolingual or a bilingual context, have pointed out that internal codas are indeed more difficult to acquire for children than branching onsets (Rose 2000; Almeida 2011).

Our aim in this study is to contribute to better knowledge of French phonological complexity in different contexts of language acquisition: monolingual, bilingual and in SLI. 110 children, aged between 5;04 and 9;00, participated in our study: 61 Typically-Developing (TD) bilingual children learning French with different L1s (either Arabic (n=25), Portuguese (n=16) or Turkish (n=20)); 20 bilingual SLI children speaking either Arabic (n=5), Portuguese (n=7) or Turkish (n=8) additionally to French; 17 monolingual French SLI children and 12 TD monolingual French children. More data from SLI children are being collected. Their elicited productions were collected using a non-word repetition task (LITMUS NWR-French), containing 71 non-words with different syllable types. For the purpose of this study, we focused on the productions of internal codas and branching onsets.

The children with SLI (monolingual and bilingual) and the TD bilinguals performed significantly lower at internal codas:

	BiSLI	MoSLI	BiTD	MoTD
Internal codas	55%	37%	77%	90%
Branching onsets	84%	66%	93%	90%

Moreover, the repair strategies used in erroneous productions also indicate that children have more difficulties with internal codas: all the cases of metathesis affecting a target internal coda result in the production of a branching onset (*e.g.* pilfu → plifu) whereas the contrary is not observed: when metathesis affects a target branching onset, it is not produced as a coda but it is rather moved in another syllable (*e.g.* fikupla → flikupa). The observed differences in (i) the rate of target-like production and (ii) the patterns of metathesis of these two structures suggest that internal codas are more complex than branching onsets for children learning French. We then examine the implications of this empirical result for theories of syllable structure.

The Great Peninsular Scandinavian Vowel Shift:
Theoretical and empirical issues

Stig Eliasson
Johannes Gutenberg University Mainz
eliasson@uni-mainz.de

Starting in the Middle Ages, Norwegian and Swedish, in contrast to Danish, underwent a sweeping chain shift of long back vowels [aː] > [ɔː] > [oː] > [uː] > [ʉː] (> Swedish [ʉ̟ː]). Several facets of this shift are typologically spectacular with numerous implications for general phonetics, phonology, and historical linguistics. The aim of the present paper is to discuss a series of controversial issues relating to the change: (1) Is the shift a case of drag or push? (2) Why does the outcome in Peninsular Scandinavian differ from that in Danish, which shows only the first phase of the process, viz. [aː] > [ɔː]? (3) Is the main trigger of the shift structural or extra-linguistic? (4) Why do we get a chain shift rather than vowel merger or diphthongization? (5) How did Peninsular Scandinavian acquire two contrasting kinds of lip-articulation, something which in Swedish resulted in the creation of the 'exotic' vowel types outrounded [yː] and inrounded [ʉ̟ː]?

The paper outlines the arguments for regarding the shift as a push chain as opposed, e.g., to Hock (1986, 1991), who views it as a drag chain (differently Labov 1994). The evidence for a push interpretation includes: (a) the intricate, quite specific ordering relationships between the component changes (Kock 1906–1929, Wessén 1968, Eliasson 1983), (b) varying geographical spread of the individual changes (Larsen 1926), (c) philological data pointing to an early date of the change [aː] > [ɔː], and (d) the rise of the typologically unparalleled contrast between outrounding and inrounding (Eliasson 1983, 2000, 2005; Ladefoged & Maddieson 1996). The reasons why Norwegian and Swedish, but not Danish, underwent the shift have long remained totally mysterious. Haugen (1970) categorically rejects a language-internal explanation, suggesting instead that "the impulses are clearly extra-linguistic", without indicating, though, what these impulses are. Eliasson (2010) draws attention to a crucial, previously entirely neglected, structure-internal difference between Danish, on the one hand, and the two other languages, most dialects of which possess a tenth vowel phoneme /ɵ/ of high frequency, yet severely stigmatized in the standard varieties. Jointly with the early, essentially pan-Nordic, change [aː] > [ɔː], the emergence of the long allophone [ɵː] triggered the chain reaction, which means that the prime cause was structural rather than extra-linguistic. The further question as to why some long back vowel pair did not simply merge or why the original high back vowel [uː] did not diphthongize is more difficult to answer. Utilizing the ingeniously organized database of Lyttkens & Wulff (1916) displaying the lexical domains of each Swedish vowel phoneme and its main allophones, the paper submits that the high lexical frequencies of all relevant vowels may have constituted a substantial obstacle to merger. Moreover, the option of diphthongizing the old [uː] was in Swedish counteracted by an absolute phonotactic requirement prescribing monophthongs in native root morphemes (the only permitted native root templates are $C_0^3VC_0^{2(3)}$ and $C_0^3VC_1^{2(3)}$vL, where lowered/raised figures indicate range of cluster complexity, v = a commonly fugitive vowel, and L = l, r, n). The odd lip-rounding contrast, finally, arose as a last expedient in severely constrained phonological surroundings. Consequently, structural factors were at work in this vowel shift to a far greater extent than envisaged by Haugen (1970). Furthermore, given its typologically unique effects, unexpected under a drag chain view, and given the specific ordering relations between the constituent individual changes, the Great Peninsular Scandinavian Vowel Shift provides an extraordinarily strong case for the existence of push chains in historical phonology.

Not-initial word-initial geminates

Luke Galea, Institute of Linguistics-Phonetics, University of Cologne
lgalea@smail.uni-koeln.de

Word-initial geminates are typologically uncommon (Kraehenmann 2011) and how they should be represented has been a bone of contention in phonology for numerous years (Davis 1994; 1999b; Hume et al. 1997; Ringen and Vago 2011). Maltese has word-initial geminates, but they have not played a large role with respect to studies of word-initial geminates. Maltese consists of three language strata: Arabic (Semitic), Italian, and English (non-Semitic). Word-initial gemination in Maltese operates as a morphophonological phenomenon which occurs in both Semitic and non-Semitic verbs and nouns. In Semitic verbs, gemination arises through the regressive assimilation of the prefix /t-/ before [+ coronal – sonorant] word-initial verbs as in (1a). On the other hand, non-Semitic verbs do not usually fit the root-and-pattern template and undergo word-initial gemination (which can affect all sounds in the language) as a process of integrating the verbs into the language (cf. Mifsud 1995).

(1) Word-initial gemination in Maltese
(1a) Semitic verbs
/t/ + /lɪbbɛs/ → /tlɪbbɛs/ 'he put something on'
/t/ + /dɐwwɐr/ → /ddɐwwɐr/ 'he was turned around'
(1b) non-Semitic verbs
/ppɐrkjɐ/ 'he parked' from English 'to park'
/ffɪrmɐ/ 'he signed' from Italian *firmare* 'to sign'

However, descriptions of Maltese (cf. Azzopardi 1981; Mifsud 1995) show that word-initial geminates may not be truly word-initial, since they are often preceded by the epenthetic vowel [ɪ]. This suggests that the language has its own way to repair a very marked structure.

In this paper, I investigate the production of word-initial geminate voiceless stops (/p, t, k/) in 12 native speakers of Maltese by measuring the duration of the geminate and the presence or absence of the epenthetic vowel before word-initial geminates. The duration of word-initial geminates is compared to that of word-medial geminates. Target words were presented in a carrier phrase *Qalilhom <target word> mitt darba* 'he told them <target word> a hundred times'. Examples of target word pairs include: /**pp**ɔntɐ/ 'he pointed' and /mɐ**pp**ɐ/ 'map'.

The results show that word-initial geminates are almost always preceded by an epenthetic vowel (96% of productions). Furthermore, the duration of word-initial geminates ($\bar{x}$=143 ms) is comparable to that of word-medial geminates ($\bar{x}$=135 ms), confirmed by a linear mixed effects model revealing that the durational difference between word-initial and word-medial geminates is not significant.

Consistent with previous claims, these measurements show that word-initial geminates in Maltese are not really word-initial as they typically require a preceding epenthetic vowel. As a result of this epenthesis, an underlyingly absent word-initial syllable surfaces. This is in line with word-initial geminates in Maltese as unsyllabifiable consonants, with epenthesis to allow their syllabification (Itô 1986). This process of syllabification in Maltese retains the geminates, however, resulting in a less marked, more typologically common position, i.e., word-medial position.

Phonetic, phonological and orthographic degemination in Dutch and German - a unified formal account

Silke Hamann (ACLC, University of Amsterdam, silke.hamann@uva.nl)

Both Dutch and German have a process of degemination of obstruents within the prosodic word (henceforth: p-word) see e.g. Booij (1995) and Ruys & Trommelen (2003) for Dutch, and Wiese (1996) for German. This is illustrated with the examples in (1), where brackets indicate p-word boundaries.

(1)	Dutch:	/(zɛt+tə)/	[zɛtə]	'to put (past)'	*zette*
		/(vud+də)/	[vudə]	'to feed (past)'	*voedde*
	German:	/(hat+tə)/	[hatə]	'to have (past)'	*hatte*
		/(liːs+st)/	[liːst]	'(you) read'	*liest*

The same process is described to occur in fast speech across p-words (ibid.), see (2):

(2)	Dutch:	/(boːt)(tɔχt)/	[boːtɔχt]	'boat tour'	*boottocht*
		/(vɪs)(sup)/	[vɪsup]	'fish soup'	*vissoep*
	German:	/(ʃɪf)(faɐt)/	[ʃɪfaɐt]	'ship travel'	*Schifffahrt*
		/(boːt)(tuːɐ)/	[boːtuːɐ]	'boat tour'	*Boottour*

Phonetic studies (Martens & Quené 1994; Jacobs et al. to appear) on Dutch, however, show that the latter process is gradual and non-neutralizing.

In the present account, we therefore propose that the obligatory process in (1) is phonological, triggered by an OCP-like constraint as in (3a), while the optional process in (2) is phonetic and due to a speech-rate dependent articulatory constraint given in (3b).

(3a) *GEMω: No geminates within prosodic words.
(3b) *[Oː]: The articulation of a long obstruent is penalized.

While (3a) applies in the mapping from underlying to surface form (in the phonological module), (3b) applies in the mapping from surface phonological form onto a phonetic realization (at the phonetics-phonology interface). We employ the bidirectional phonetics-phonology model by Boersma (2007) to formalize this principal distinction within OT. In this model, high-ranked phonological constraints such as (3a) also restrict the output of the perception process (the phonological surface form), and can therefore account for the fact that e.g. the Italian phonetic form [latːe] 'latte' is perceived as containing a singleton /latə/ both by German and Dutch native speakers.

In the present account, we furthermore propose that (3a) can also be employed to restrict the reading process, i.e. the mapping from an orthographic form onto a surface phonological form. Here, phonological constraints like (3a) interact with orthographic mapping constraints such as the one given in (4):

(4) $<\beta_i\beta_i>$/C/: The mapping of two identical consonantal graphemes onto a single surface consonant is penalized.

In our analysis we show that this constraint has to be outranked by (3a), and is supplemented by an undoing of the degemination process in the mapping from surface to underlying form, in order to account for the cases in (1) and (2) and also for monomorphemic words such as German *Matte* 'mat' with an underlying singleton.

The proposed formalization of the reading process as an interaction of phonological with orthographic constraints is preferable to earlier accounts such as Neef (2012), which reduplicated phonological restrictions in their orthographic mappings.

Rethinking the *Core-Periphery* Model: Lexical Evidence from Japanese

Ryan Hearn – Cornell University – rdh238@cornell.edu

Ito and Mester (1995, henceforth I&M) developed the *Core-Periphery* model to account for apparent behavioral differences between strata in the Japanese lexicon. By assigning lexical items to strata with distinct rankings of FAITH, I&M attempt to account for these alternations with one overall ranking of markedness constraints. In this paper, however, I argue that not even this reranking of FAITH is necessary, and that all Japanese lexical data, including the numerous exceptions to their proposed strata, can be evaluated using a single ranking of constraints, if we keep in mind the reality of lexical change.

I suggest that the hierarchical nature of lexical stratification is due to the constraint reranking process that results in long-term phonological change, and that the phonotactic similarities within each stratum are a reflection of the constraint ranking active during the period when these forms were initially lexically stored. Once we accept that underlying forms change from generation to generation, and allow them to be updated accordingly, we realize that alternations that have posed problems for OT when attributed to phonology can be easily dealt with through the lexical representation instead.

I&M state that over time lexical items move toward the core of the lexicon and become more constrained, but I present three sources of data which show the opposite. I argue that the paradigm of the suffix *pun* 'minute' has levelled in casual speech from *i-ppun* 'one minute', *ni-ϕun* 'two minutes', *san-bun* 'three minutes' to *ip-pun*, *ni-pun*, *san-pun*, in violation of two of I&M's constraints that should apply to *pun*'s stratum: "postnasal obstruents must be voiced" (NO-NT), and "no singleton-p" (NO-P). Rice (1997) also provides evidence of core lexical items violating I&M's NO-NT constraint, including *anta* 'you', formed via syncope from *anata*. I&M are aware of these exceptions, but refer to them as "undoubtedly native, but peripheral" (830). Additionally, I&M use a constraint preventing voiced obstruent geminates to separate their two most peripheral strata, Assimilated and Unassimilated loans. According to Crawford (2009), however, there are only ten forms with relevant devoiced geminates, and at least five different adaptation mechanisms for loans with final voiced consonants, making this constraint an odd choice to differentiate these strata.

I argue instead that FAITH outranks both NO-P and NO-NT outright in Modern Japanese, and that all modern surface alternations seemingly due to these markedness constraints are instead lexically stored. In the Old Japanese period, however, the opposite ranking held; both outranked FAITH, effectively eliminating from the language any singleton-p and NT clusters that existed previously. Later in the Sino-Japanese period, under the influence of massive numbers of borrowings from Chinese, the NO-NT constraint was reranked below FAITH, allowing new words with NT clusters to enter the language. Of course, since NT clusters had been eliminated from all Old Japanese words previously, they were now missing from the entire surviving Old Japanese stratum in Middle Japanese. Crucially, however, this is now due to lexical transmission and acquisition and no longer due to constraint interaction. This pattern then repeated with the influx of Western borrowings in the 1600s, reranking NO-P below FAITH, but leaving the entire Middle Japanese stratum singleton p-less, again for lexical reasons. Since these restrictions were no longer phonological, the Old Japanese word *anata* was free to syncopate to *anta*, and the paradigm of *pun* was able to level analogically.

This paper advocates rethinking the boundary between phonology and the lexicon, reconsidering much of the reliance upon strata within OT, and stressing the importance of lexical representation in the input to OT evaluation.

Pre-aspiration, sonorant devoicing, and the sonority hierarchy

Míša Hejná
Newcastle University
misprdlina@gmail.com

This study illustrates two points. Firstly, pre-aspiration and sonorant devoicing, at least in various accents of Welsh English, are the same phenomenon affecting different types of segments. Secondly, the presence of sonorant devoicing (as in *li**nt**er* [lɪn̥tˢə]) implies the presence of pre-aspiration in the same language (as in *li**tt**er* [lɪʰtˢə]), but not the other way round (i.e. the presence of pre-aspiration does not imply the presence of sonorant devoicing). More precisely, the phenomenon affects different types of segments that occur before voiceless obstruents in the following order: V > liquids > nasals (*litter* > *Hilton* > *linter*). This pattern corresponds to the sonority hierarchy, including the fact that low vowels are affected more frequently than high vowels.

48 speakers, who were born and raised in various parts of Wales, provided 2483 tokens n total relevant for the analyses of pre-aspiration and sonorant devoicing within every individual. The segmental and prosodic structure of all the tokens taken into consideration is that of ˈCVPV (*litter*), ˈCV:PV (*naughty*), ˈCVCPV (*linter*), and ˈCV:CPV (*haunter*). The data collection was done via wordlist.

Pre-aspiration was defined as a period of voiceless (primarily glottal) friction and sonorant devoicing was similarly identified based on the presence of voiceless (primarily glottal, but often also oral) friction. The voiced interval of breathiness, which is often considered part of pre-aspiration (e.g. Morris 2010; van Dommelen 1999, 2000; van Dommelen, Holm & Koreman 2011) was not included in the analyses because breathiness is problematic to analyse acoustically in nasals and the comparisons across the different types of segments would therefore also be problematic. None of the vowels or consonantal sonorants are fully devoiced and the voiceless friction is unsurprisingly found closer to the source of this voicelessness (i.e. the plosive).

The paper will discuss the consequences of devoicing for the preservation of the place contrast for the sonorants (Silverman 1996: 364; Shosted 2008) as well as the plosives (Mann 1980: 411; Mann 1986) and compare the present results with those found in Sienese Italian, in which sonorant devoicing (tested on /l/) is more frequent (85%) than pre-aspiration (48%; Stevens & Hajek 2004).

Nasal gemination in Standard Mandarin (SM) loanwords: Corpus vs. experimental results

Ho-Hsin Huang (huanghoh@msu.edu), Michigan State University

This study investigates how English intervocalic [n] is adapted into SM loanwords both in the corpus and in perceptual similarity adaptation data from SM monolinguals and SM-English bilinguals. The findings show that (i) the bilingual and monolingual experimental strategies for nonce word adaptations are similar to the observed patterns in the corpus; (ii) both monolingual and bilingual groups depend on the perceptual cues to form loanwords. I account for this process by showing how perception and native SM phonology interact.

Corpus Data & generalizations. The exiting corpus presents two main factors that trigger nasal gemination in SM loanwords: (i) the prenasal vowel should be lax (e.g *ˈDaniel* →[ta**n.n**iːəɹ]; **Bruno*→[puːluː**n**^{w}o]), and (ii) the prenasal vowel should bear the primary stress (e.g. ˈ*Benny*→[pa**n.n**iː],**De*ˈ*nise*→[tiː**n**iːsɻː]). However, when the stressed vowel is a low back [ɑ], nasal gemiantion does not occur (e.g. *O*ˈ*bama*→[au.baːmaː]).

Experiments & results. To test the extent to which the generalizations identified in the corpus are based on perception, a nonce word perceptual similarity adaptation task was run on 33 Taiwanese SM monolinguals and 24 SM-English bilinguals. The test items have an English [n] (a) with the prenasal vowel lax/tense or (b) with the prenasal vowel stressed/unstressed. The participants listened to two possible SM adapted forms followed by an English input (e.g. SM[ban.niː], SM[baː.niː], ENG[ˈbæni]), and chose the most perceptually similar SM output to the English input. I compared the preference of bilingual and monolingual speakers for nasal gemination separately for both vowel type and stress conditions. T-tests of the vowel type condition reveal that the preference for nasal gemination is significantly higher in both monolingual ($t(32)=13.99$, $p<0.001$) and bilingual ($t(23)=9.19$, $p<0.001$) groups when the English prenasal vowel is lax. When the English prenasal vowel is [ɑ], the nasal germination preference in monolingual ($t(32)=7.76$, $p<0.001$) and bilingual ($t(23)=7.62$, $p<0.001$) is significantly lower than other lax vowels. T-tests of the stress condition also reveal that the preference for nasal gemination is significantly higher in both monolingual ($t(32)=10.13$, $p<0.001$) and bilingual ($t(23)=7.58$, $p<0.001$) groups when the unstressed prenasal vowel is lax.

Discussion & conclusion. Nasal gemination is an unnecessary repair (Kang 2011) since SM syllable structure allows a single nasal for a direct match. I propose that nasal gemination is needed because a VN rime in SM better matches the duration and nasality of the prenasal lax vowel in English (Hayes 2009, Solé 1992) and also satisfies the SM μμ-syllable constraint (Duanmu 2007). In both the vowel type and stress conditions, English tense vowels and the phonetically long vowel [ɑ] are perceived as long and ready for a direct syllable mapping in SM. Nasal gemination cannot apply since a V:N rime is illicit in SM. Unlike the patterns in the corpus data, speakers in both groups chose nasal gemination when the prenasal vowel is unstressed and lax. However, the English input recordings show clear nasalization on the unstressed prenasal lax vowels. This indicates that both groups of speakers rely on perception cues to match vowel duration and nasality by means of nasal gemination. Under the stress type condition, neither group shows a nasal gemination pattern as consistent as the corpus data when the prenasal stressed vowel is lax, but both groups do show a smaller nasal gemination preference when the prenasal stressed vowel is tense. Overall, the experimental results seem to suggest that input vowel duration provides a crucial perceptual cue for the nasal gemination process.

In conclusion, I propose that nasal gemination is guided both by perceptual assimilation, which promotes similarity between the perceived input and the selected output, and by the native phonology of SM, but online adaptation tasks induce more reliance on perceptual cues.

On Accentual Preference of English Intransitive Sentences

Yujing Huang, Harvard University (yujinghuang@fas.harvard.edu)

It has been observed that the prosody of English intransitive sentences can vary in a broad focus context, i.e., some sentences can have their sentential focus realized on the subject only while others need to put accents on both the subject and the verb. Some linguistic analyses explain the variation by the difference in the unaccusaitivity of the intransitive verbs. The proposal is: because the subject of an unaccusative verb originates from the object position and leaves a trace at the base position, realizing the focus on the subject licenses F-marking of trace (i.e., the interval argument of the verb) and therefore licenses F-marking the verb phrase (Selkirk, 1995). This proposal has been revised under the Minimalist framework with phase theory (Kratzer and Selkirk, 2007). However, whether the variation in focus realization is due to unaccusativity is controversial. Although some empirical research agreed with such an observation (Hoskins 1996, Irwin 2011), some studies failed to find a difference driven by unaccusativity (Hirsch and Wagner, 2011). These studies differ in their measurements, experimental tasks and the control of potential confounds, and therefore make the inconsistency hard to interpret.

In the current study, 34 test sentences were elicited by the question "what's up" for a broad focus. To control for the confounds, all the unaccusative verbs were paired with unergative verbs and each pair had the same human subject (e.g. manager). The frequency of the verbs were matched and the predictability between the subject and the verb was controlled following the same method of Verhoeven and Kügler (2014). The F0 and duration of the subjects and verbs were measured. The F0s were z-score normalized for each participant. In addition to the measurements of the previous studies, this study also measured the boundary between the subject and the verb. The rationale is that, if a sentence has subject only accent, there should be no prosodic boundary between the subject and the verb. This predicts a weaker final lengthening between the subject and the verb in unaccusaitve sentences (measured as the duration of the final syllable of the subject).

No significant difference is found between the unaccusative condition and the unergative condition. Different mixed effect models with verb type as the main effect, sentence, participants, presentational order and number of syllable of the verb as random effect, F0 of the subject, F0 of the verb, duration of the subject, final lengthening as dependent variable respectively shows that there is no main effect of verb type on peak F0 of the subject ($p = 0.64$), on peak F0 of the verb ($p = 0.67$), on duration of stressed syllable of the subject ($p= 0.88$) or on the boundary ($p = 0.62$).

Because unaccusative sentences and unergative sentences can have accents on both the subject and the verb and there is no predicted difference in this prosodic pattern between the two verb types, one can argue that the difference in focus pattern between unaccusative sentences and unergative sentences is washed off by averaging. If this is true, we expect that unaccusative sentences can have two prosodic pattern and unergative sentences can only have predicate prominence (Selkirk, 1995). The TOBI analysis shows that this prediction is not borne out. 39% of the unaccusative sentences have subject only prominence while 28% of the unergative sentences have subject only prominence (Fisher's Exact Test $p = 0.29$). This shows that the unaccusative and unergative sentences both have predicate prominence as well as subject prominence, and there is no systematic difference in the prosodic structure between these two types of sentences. This result cannot be explained by the existing focus projection theory and leaves us to explain where the variation comes from.

Phonetically real, but phonologically impossible: The lack of RVA in Italian

Bálint Huszthy, Péter Pázmány Catholic University, huszthy.balint@gmail.com

The basic hypothesis of the present survey is the claim that Italian is a voice language, in which *regressive voice assimilation* (hereafter RVA) is completely missing. The study puts forth four arguments to demonstrate the lack of RVA in the synchronic phonology of Italian, whereas the data are based on voice recordings made with 79 Italian informants.

First of all, the corpus reveals that the voicing of /s/ before voiced consonants – formerly considered in Italian linguistics as voicing assimilation (cf. Nespor 1993; Schmid 1999; Krämer 2009) – vacillates in the pronunciation of recent loanwords, e.g. (a) *iceberg* [ˈajsbergə], (b) *back*[s/z]*lash*, (c) *kri*[s/z/ʃ]*na*, (d) *kala*[ʃ/ʒ]*nikov*, etc. If /s/ does not necessarily get voiced before voiced obstruents, as in (a), or it vacillates, as in (b-d), its potential voicing cannot be considered in synchrony an automatic postlexical process, like RVA.

Secondly, /s/ never undergoes voicing in sandhi position, unlike in the case of traditional RVA, e.g. (e) *Champion*[s] *League*, (f) *silen*[s] *drive*, (g) *Pier*[s] *Bro*[z]*nan*; *autobu*[s] *bianco* 'white bus' (cf. Bertinetto–Loporcaro 2005), etc. Thirdly, affricates which hold the [sibilant] feature, usually do not get voiced before voiced consonants, e.g. (h) *u*[t͡s/z]*beco* 'Uzbec', (i) *sami*[t͡s/z]*dat*, (j) *gu*[t͡s/z]*la* 'gusle', etc.

Finally, other obstruent clusters in Italian always tend to conserve their original voice value. The most common repair strategy to achieve this is schwa-epenthesis, but several examples also are found without a schwa and with the entire conservation of the consonants' voice value, e.g. (k) *vo*[dk]*a*, (l) *M*[ekd]*onald's*, (m) *ou*[td]*oor*, (n) *e*[kd͡z]*ema* 'eczema', etc.

All the same, from the point of view of *laryngeal realism* (Iverson–Salmons 1995; Honeybone 2002, 2005; Cyran 2014), Italian is definitely a voice language, as well as the other Romance languages, since voiced obstruents can carry voice even in non-intersonorant position, and voice is phonologically distinctive. Moreover, aspiration is not systematic in Italian (except for Tuscan varieties), and it is never distinctive (cf. Marotta 2008).

Despite the solid phonetic evidence for the lack of RVA, this study aims to claim that the current laryngeal status of Italian is phonologically impossible: indeed the adjacency of voiced and voiceless obstruents is ill-formed from a phonological approach. The spectrograms attest that the informants try to resolve this situation in different ways, but the only successful solution is the schwa-epenthesis (e.g. *vod*[ə]*ka*, *out*[ə]*door* etc.). Apart from that, speakers seemingly turn down the voice measure of the voiced consonant at the edge of the voiceless one, or they leave a minuscule silence (about 1,2 ms) between the two elements, or they just extend the release phase of the first consonant before the second one. Apparently, the informants are somehow aware that what they phonetically do, is impossible from the point of view of phonology, and they involuntarily attempt to avoid this situation.

However, the most interesting part of the corpus concerns the pronunciation of five Italian-Hungarian bilingual (or near bilingual) informants. Hungarian is clearly a voice-language, and every obstruent cluster undergoes RVA (e.g. in the Hungarian pronunciation of the loanwords *a*, *f-i*, *k-n* among the above examples). These five informants apply RVA almost at every occurrence (except in the case of schwa-epenthesis), but only they. According to the other informants' opinion, these bilingual speakers do not have any foreign accent in their Italian speech (but they speak Hungarian with an Italian accent).

I assume that once Italians acquire RVA through a strong contact with another voice language, they automatically add it to their L1 phonological store, and they keep using it in L1 pronunciation. To sum up, the current lack of RVA in Italian is a provisional synchronic situation; and if RVA is phonetically not obligatory in voice languages, phonologically it is.

Local faithfulness constraints over correspondence structures

Adam Jardine, University of Delaware - ajardine@udel.edu

What is the nature of phonological FAITHFULNESS? The current work defines a strictly *local* notion of FAITHFULNESS through *banned substructure constraints* over *correspondence structures* that explicitly represent the correspondence between the input and the output in a phonological transformation. By making these constraints *language-specific* and *inviolable*, we can study the local nature of FAITHFULNESS in phonological generalizations without recourse to optimization, which is known to derive non-local interactions from local constraints.

Most mainstream theories of phonological transformations specifies a set of input structures paired with output structures (Chomsky & Halle, 1968; Prince & Smolensky 1993, 2004; McCarthy 2000, 2010). For example, consider the intervocalic voicing generalization "obstruents are voiced intervocalically'; a rule is given in (1). Independent of whether one analyzes it with rules or constraints, in this generalization /pa/ would be paired with the faithful [pa], but /apa/ would be paired with [aba]. An additional notion is that units in the input *correspond* to specific units in the output (Kager, 1999; McCarthy and Prince, 1995). Thus, in /pa/→[pa], the input /p/ corresponds to the output [p], but in /apa/→[apa] /p/ corresponds to [b], as in (2).

(1) $[-\text{sonorant}] \rightarrow [+\text{voice}]$ / V __ V

(2) a. p a / p a b. a p a / a b a

(3) G^a_a= a / a G^p_p= p / p G^p_b= p / b

(4) a. * a p a / a p a b. * p a / b a

We can define constraints over correspondence structures like in (2) (Potts and Pullum, 2002). First, to define a universal set of possible correspondence structures (e.g., GEN in OT), we can use the *concatenation* of a finite set of individual primitives (Jardine and Heinz, 2015). Here, the primitives are possible phoneme/phone correspondences, as given in (3); the /apa/→[aba] structure in (2b) is obtained by $G^a_a \cdot G^p_b \cdot G^a_a$ (where $\cdot$ indicates concatenation).

This concatenation can also generate the correspondences in (4), which are invalid for intervocalic voicing. Thus, there must be some way of specifying a language-specific set of valid correspondences. In OT, this is done through finding the optimal correspondence structure over a series of violable constraints, but this can generate unattested, *non*-local patterns from local constraints (Gerdemann and Hulden, 2012; Heinz and Lai, 2013; Wilson, 2003). To ensure locality, we can turn to a computational notion of local constraints based on banned substructures (Rogers et al., 2013; Jardine and Heinz, in press). For example, a set of such structures for the intervocalic voicing generalization are given in (5), specifying an intervocalic /p/ that has surfaced faithfully and an initial /p/ that has surfaced unfaithfully, respectively. We can define the set of correspondences in the intervocalic voicing generalization by banning these substructures, as the invalid correspondences in (4) contain them, as highlighted in bold in (6).

(5) a. * p / a p a b. * # p / b

(6) a. * a **p** a / **a p a** b. * **p** a / **b** a

In this way, whether or not FAITHFULNESS is adhered to in a particular correspondence is determined entirely by neighboring units in the substructure—it is in this sense that it is *local*. This notion of locality can be recruited for learning (García et al., 1990; Jardine and Heinz, 2016), and it can also be extended to non-local transformations through different interpretations of the substructures (Heinz, 2010; Rogers et al., 2013).

Effects of Neighborhood Density and Spelling on Adult Word Learning

Skott Jones (sejones@ithaca.edu), Taryn Carlson, and Cassandra Stolting
Ithaca College, Ithaca, NY USA

During their first several years of life, children gradually become attuned to the language that they are learning. Some of their remarkable accomplishments include hearing the phonological representations of novel words, and quickly incorporating them into their lexicon. Beyond this critical period, however, individuals are tasked with continuous learning of novel words in order to successfully maneuver through academic and professional environments. Neighborhood density (ND), a variable proposed to operationalize phonological similarity, refers to the relationship between a word and the number of similar-sounding words (i.e., dubbed *neighbors*) in the lexicon. Words are considered neighbors if they are identical in all aspects except for the addition, deletion, or substitution of a phoneme (Luce & Pisoni, 1998; Vitevitch & Luce, 1998). For example, *cat* and *cut* are neighbors. However, *cat* and *coach* are not neighbors because they differ in more than one sound. Words with many neighbors are considered to be *dense*, and words with few neighbors are termed *sparse* (Blumenfeld & Marian, 2006).

Despite several studies on effects of ND on word learning in children, few experiments have examined how ND might influence word learning in adults. Previous studies have reported conflicting results (dense = sparse, Freedman, 2015; dense > sparse, Storkel, Armbruster, & Hogan, 2004), prompting further research in the area of ND and adult word learning. Additionally, there has been little to no investigation of how the orthographic representation (i.e., spelling) of a word may interact with influences of ND.

The purpose of the current study was to determine if college-aged adults acquire dense or sparse words to a greater degree, particularly during a lecture-style presentation, and also whether a simultaneous presentation of an orthographic representation might impact the results. Across two experiments, 50 adult participants were exposed to a lecture about a fictitious culture. During the lecture, eight novel words (4 dense, 4 sparse) were paired with novel visual stimuli. In the first experiment, 25 participants were exposed to solely the phonological form of the words, while in the second experiment 25 participants were also exposed to the spelling of the presented stimuli. Participants were asked to recall these novel words across different levels of exposure: following one exposure, four exposures, and a 10-minute distractor task.

Findings indicated that when participants were exposed solely to the phonological forms of words, segmental detail was most accurate on dense words. This varied by phase of testing, however. Dense words were learned more accurately than sparse words following four exposures, and again after a 10-minute distractor task. Regarding the influence of spelling, participants learned words best when the orthographic form of a word was provided. Interestingly, ND did not impact learning regardless of phase of testing in this context. That is, no significant differences were found between sparse or dense words in the binary or segmental analyses when orthographic forms were provided.

In summary, when phonological forms were solely provided to participants, words with more neighbors facilitated word learning. This could be due to increased lexical activation of these forms from similar-sounding words. And yet, participants learned words best when both orthographic and phonological forms were presented. One possible explanation for this finding is when provided with orthographic forms, participants relied more heavily on visual memory to learn the novel words. Alternatively, perhaps participants learned words best when multiple cues were provided (i.e., auditory and visual).

Deconstructing [continuant] in the coda – a typological survey

Martin Krämer
UiT The Arctic University of Norway

Draga Zec
Cornell University

We investigate the organization of coda segmentism, focusing on the role of manner of articulation features. According to an influential proposal (Clements 1990), the inventory of segments in the coda universally observes the sonority hierarchy, with less sonorous segments in the coda of a CVC syllable implying the presence of the more sonorous ones. With the sonority scale L>N>F>P (Liquids > Nasals > Fricatives > Plosives), the predicted coda inventories include solely {L}, {L,N}, {L,N,F} and {L,N,F,P}. However, our database of 204 geographically and genetically spread languages paints a different picture. First, we document coda inventories such as {N}, {P}, {N,P}, {L,N,P} or {N,F}, which do not conform to the sonority based predictions. Second, one of the predicted inventories, {L}, does not figure in our database, even though it should be particularly privileged from the perspective of sonority. In sum, sonority is not a reliable predictor of manner of articulation classes found in codas cross-linguistically.

As an alternative, we propose to explore how the feature [continuant], classified in the SPE as a manner feature (Chomsky & Halle 1968), fares as a predictor of the set of attested coda inventories. This feature, however, has been defined in the literature in at least three ways. In addition to the SPE definition in (b) below, we also find (c), which refers to a specific region in the mouth, and (a), which incorporates the entire vocal tract.

(a) Unimpeded air stream anywhere in the vocal tract (Mielke 2005)
(b) Unimpeded air stream in the "vowel tract", that is, in the mouth: "(Primary constriction in) the vowel tract is not narrowed to the point where the air flow past the constriction is blocked." (SPE, p.317)
(c) Unimpeded air stream in the midsaggital region of the oral tract: "[V]ocal tract configuration [allows] the airstream to flow through the midsaggital region of the oral tract." (Halle & Clements 1983)

None of the three definitions alone can capture the set of attested inventories. Instead, we propose that all three definitions are active cross-linguistically, and that individual languages select one, which then shapes the language's coda inventory, and is also consistent with the workings of its phonological system. That is, in the spirit of Mielke's (2004, 2005) proposal, we allow for features to have language specific interpretations. However, we also assume that a feature such as [continuant] has a universal status, and that languages can only differ in the size of the articulatory domain that they select. Thus, /l/ is [+continuant] according to (a) and (b), and [-continuant] according to (c), yielding both {L,F} and {L,N,P} as natural classes. Nasals are [+continuant] according to (a), and [-continuant] according to (b), (c), giving rise to both {N,F} and {N,P}. Assuming a constraint *[+continuant]/Coda yields coda inventories with only {P} under definition (a), {P,N} under definition (b) and {P,N,L} under definition (c). Notably, the {L,P} set, excluded on this ground, is not an attested case.

On a language-specific basis, liquids and nasals can be underspecified for [continuant], which makes them immune to coda neutralization. Thus, coda neutralization either bans [+continuant] or any specification of [±continuant] from the coda (*[αcontinuant]/coda). There is an implicational relation between N and L: If Ls are underspecified, so are Ns, but not vice versa, yielding {N} and {L,N}, but no {L}. Such underspecification is supported by the phonetic and phonological ambiguity of nasals and liquids that led to the postulation of the different definitions of [continuant] and explains the high frequency of the sonorants in coda inventories as well as the absence of liquid-only inventories.

The typology of coda inventories can be understood as variations on the feature [±continuant], if i) the feature has a universal semantic core that is subject to language-specific interpretation with regard to its scope in the vocal tract, ii) all sonorants can be underspecified or iii) only the nasals can be underspecified for [±continuant].

Stress retraction and its consequences for Gaeilge Chorca Dhuibhne phonology

Anton Kukhto
Lomonosov Moscow State University
anton.kukhto@gmail.com

The present paper focuses on one aspect of the prosodic phonology of Gaeilge Chorca Dhuibhne (Dingle peninsula, Co. Kerry, Ireland), one of the Munster dialects of Modern Irish. In it the mechanism behind the so-called lexical stress retraction is analysed and the consequences it yields for the whole phonological system are shown.

By stress retraction I hereby understand the leftward "movement" of primary lexical stress in a given word under the influence of this word's immediately following context in a phrase, e.g. *cailín* [kalʲˈiːn̪ʲ] 'girl', but *cailín óg* [ˌkalʲiːˈn̪ʲoːg] 'young lass'. The literature on the subject does not make it clear what the conditions under which the stress might be retracted are and even what types of words can undergo this "transformation", for example, Bennett (to appear) notes that retraction can occur in words "which have an [LH] weight profile"[1], and Ó Sé (2000: 52–3) adduces some examples where stress retraction is observed. One further remark to be made is that lexical stress assignment in Gaeilge Chorca Dhuibhne depends on the syllabic composition and the moraic structure of words (see Ó Sé 2000, Iosad 2013, Torres-Tamarit & Hermans 2016 for details).

I claim that the aforementioned factors also condition stress retraction in Gaeilge Chorca Dhuibhne, i.e. the requirements for rhythmic improvements have to be formulated in terms of moraic structure. Taking into account the available sources (see above) and a corpus of my own field recordings, I show that a situation in which there is pressure on the stress to move forward arises within an intonational phrase in those occasions when the distance between the two underlying primary stresses is less than or equal to two morae (≤2). This rule helps to acknowledge the fact that cases like *múinteoir* [muːn̪ʲˈt̪ʲoːɾʲ] 'teacher' ~ *múinteoirí scoile* [ˌmuːn̪ʲt̪ʲoːɾʲiː ˈskɪlʲə] 'school teachers' (with a bimoraic syllable in between) and *macalla* [məˈkalə] 'echo' ~ *macalla na haille* [ˌmakələ nə ˈhalʲə] 'echo in the cliffs' (with two monomoraic syllables in between) behave equally, while cases like *amhránaí* [aʊˈɾɑːn̪iː] 'singer' ~ *amhránaí as Albain* *[ˌaʊɾɑːn̪iː əsˈ aləbən̪ʲ] 'singer from Scotland' (with three morae in between) are not found in the data[2], despite the fact there is no metrical restriction on such a structure, cf. the first example in this paragraph.

It is also of interest that there exist types of words that disallow stress retraction under rhythmically satisfying conditions, the largest group of them being words with the element /ax/ in the second stressed syllable (which is an exceptional kind of behaviour, for no other second light syllable in a LL sequence can be stressed, see Ó Sé 2000). Given the fact that in other cases retraction can recover otherwise unstressed vowels of the first syllable, we are then led to believe that stress retraction sheds light on vowel reduction in Gaeilge Chorca Dhuibhne. Contra Bennett (to appear), who believes reduction to be a unitary postlexical process, I propose a "split" model of vowel reduction, which assumes that a substantial proportion of unstressed vowels indeed undergo a postlexical process of reduction, but others are reduced underlyingly and are stored in the lexicon. Not only these vowels react to stress retraction in an unusual manner, but also they demonstrate distinctive phonetic behaviour as far as their lengths and spectra are concerned.

---

[1] Bimoraic syllables with long vowels are H(eavy) and monomoraic, with short ones, are L(ight).

[2] Needless to say, the absence of something from the data doesn't prove its definitive impossibility, but it is worth noting that speakers also report examples of the like to sound unnatural.

Laryngeal contrast in Qatari Arabic: Aspiration, voice, or both?

Vladimir Kulikov, Qatar University, vkulikov@qu.edu.qa

Typological studies of voicing (e.g. Iverson & Salmons, 1995) show that languages typically contrast voiceless unaspirated stops (*p, t, k*) either with (pre)voiced stops (*b, d, g*) or with voiceless aspirated stops (p^h, t^h, k^h). The type of contrast in a language can be identified by changes of voice onset time (VOT) in response to speaking rate manipulation. These changes are asymmetrical: in slower speech, negative VOT (prevoicing) in voiced stops and long-lag VOT (aspiration) in voiceless aspirated stops increase, but short-lag VOT in unaspirated stops does not (Kessinger & Blumstein, 1997).

Beckman, Helgason, McMurray & Ringen (2011) argue that this asymmetry can be linked with phonological specification of stops. Under this theory rate only affects VOT in prevoiced stops, specified with privative [voice], and in voiceless aspirated stops, specified with privative [spread glottis]. VOT in phonologically unspecified voiceless unaspirated stops is not affected by rate. Beckman et al. (2011) also claim that some languages can have an over-specified phonological contrast between prevoiced and voiceless aspirated stops and support this claim by the finding that in Swedish both prevoicing and aspiration increase as speech slows. Yet, Swedish is the only language in which this pattern has been found so far.

The current study reveals a similar phonological pattern in the vernacular Arabic dialect of Qatar. Unlike other dialects of Arabic that have a contrast between voiceless aspirated and unaspirated stops, e.g. Saudi Arabic (Flege & Port, 1981; AlDahri, 2013) and Jordanian Arabic (Mitleb, 2009), or a contrast between voiceless unaspirated and prevoiced stops, e.g. Lebanese Arabic (Yeni-Komshian, Caramazza & Preston, 1977; Khattab, 2000), Qatari Arabic shows evidence for both prevoicing and aspiration.

To determine the type of a laryngeal contrast in Qatari Arabic, acoustic parameters of initial stops and rate effects on VOT were investigated. The data were collected from eight native speakers, who read words (n = 50) with voiced (*b, d, g*) and voiceless (*t, k*) stops in a carrier phrase at two rates (slow, fast). The results showed that 77% of voiced stops were prevoiced (Mean VOT = –69 ms), and all voiceless stops were aspirated (Mean VOT = 55 ms). Measurement of four additional cues to voicing (SCG of burst, f0, F1, and duration of the following vowel) showed significant differences ($p < 0.01$) between the two laryngeal categories in all cues.

Importantly, the effect of rate on VOT was highly significant: VOT in both voiced and voiceless stops increased in slow speech (MD_{voiced} = 23 ms, $p < 0.0001$; $MD_{voiceless}$ = 11 ms, $p < 0.0001$). In addition, the results showed a significant relation between VOT and word duration used as a proxy of speaking rate (Voiced: $R^2_{change} = .141$, $p < .0001$; Voiceless: $R^2_{change} = .073$, $p < .0001$).

The results suggest that the VOT pattern found in voiced Qatari Arabic stops *b, d, g* is consistent with prevoicing in Dutch or Swedish. The pattern found in voiceless stops *t, k* is consistent with aspiration in German or Swedish. Similar to Swedish, both prevoicing and aspiration in Qatari Arabic stops change in response to speaking rate, but the magnitude of these changes is smaller. The findings provide further empirical support for naturalness of an over-specified phonological contrast found in Swedish.

The results also reveal that the patterns were somewhat different in the two speech communities in Qatar. In our pool, four Bedouin (i.e. 'nomadic', or rural) speakers produced voiceless stops with significantly shorter aspiration than four Hadar (i.e. 'sedentary', or urban) speakers did ($M_{Bedouin}$ = 47 ms, M_{Hadar} = 63 ms; $p < 0.01$). Mean duration of prevoicing did not differ significantly between the two groups. Shorter aspiration in the Bedouin community is consistent with the traditional Arabic pattern reported by Flege & Port (1981) and Mitleb (2009), but longer aspiration in the Hadar community is more consistent with aspiration in languages like English or German.

Is Ukrainian a bidirectional stress system?

Beata Łukaszewicz
University of Warsaw
b.lukaszewicz@uw.edu.pl

Janina Mołczanow
University of Warsaw
jmolczanow@uw.edu.pl

This paper investigates phonetic underpinnings of lexical and rhythmic stress in Ukrainian, a language which is potentially an example of an intricate bidirectional stress system, as yet not discussed in the literature. The study is part of a larger project on subsidiary stress in Ukrainian. The data are interesting for two reasons. First, they run afoul of a recent hypothesis that bidirectional stress systems are non-existent (Newlin-Łukowicz 2012). Second, they reveal a complex interplay between the main stress and subsidiary stress, which is relevant from the point of view of a metrical theory separating lexical accent from rhythmic beats (van der Hulst 2014). According to traditional descriptions, Ukrainian has lexical stress and rhythmic stresses radiating from the word edges towards the syllable carrying primary stress, e.g. *ˌpereˈpysaˌno* 'rewritten', *ˌveloˌsypeˈdyst* 'cyclist', *ˈbačyˌtymeˌte* 'you will see' (Ziłyński 1932, Bilodid 1969). However, stress patterns in longer words, exhibiting both iteration and directionality, have been largely ignored. The main phonetic correlate of primary stress in Ukrainian is duration (Toc'ka 1981). Secondary stress has been posited based on impressionistic evidence. Traditional grammars are inconsistent in their descriptions of acoustic manifestations of lower stress levels, and the possibility that the rhythmic beats could also be cued by duration has not been confirmed so far.

The present paper reports on an acoustic study of the Ukrainian stress patterns based on 30 single-root words, collected from 16 monolingual native speakers of Ukrainian, living in the area of Drohobych (western Ukraine), in a word-list reading task. Only the rightward iteration of the rhythmic stress was considered. Since the study focuses on stress iteration and directionality, the data set predominantly consisted of words in which the initial syllable and the syllable with lexical stress were separated by at least three positions, e.g. *ve.lo.sy.peˈdyst* 'cyclist', *fe.ty.shy.zuˈva.ty* 'fetishize', *te.le.fo.niˈza.ci.ja* 'telephonization'. Whenever possible, counterparts with a potential lapse on the pretonic syllable were added, e.g. *ve.lo.syˈpe.dnyj* 'bycycle'. Duration measurements were conducted in six different syllable positions: initial, second, iterative, fourth, pretonic, and tonic. Dependent measures were used: the length of each syllable was expressed in relation to the mean syllable duration within a word (cf. Beckman 1986, Levi 2005). The differences between syllable positions were statistically significant ($F(5, 437.069) = 239.04$, $p < 0.001$). Pairwise Games-Howell post hoc comparisons confirmed that the syllable carrying lexical stress was significantly longer than all other syllables in a word, reaching about 1.5 the average length. The second syllable turned out the shortest (about 0.7 the average). Interestingly, the initial, iterative, fourth and pretonic positions did not emerge as significantly different. The statistical results showed that the fourth syllable was only minimally shorter than the preceding iterative syllable. Thus, a clear indication of an alternating subsidiary stress pattern was detected only in the initial-second and second-iterative sequences.

The duration measurements also revealed that the syllable immediately preceding main stress was unexpectedly long, the fact not reported in the traditional descriptions of Ukrainian. Following Bethin's (2006) account of pre-tonic lengthening in East-Slavic dialects, we hypothesize that the increased duration of pretonic syllables in Ukrainian is not rhythmically conditioned, but might be caused by a stress-tone interaction in which lexical tone is associated with a syllable immediately preceding main stress. However, the lengthening effect is very subtle in comparison with that reported in the East-Slavic dialects and there are no F_0 measurements at present to (dis)confirm the assumption of the tone-induced duration in Ukrainian. Therefore, we leave this matter for further research.

Parallel Patterns Between Language Games and Serial Music

Sara Mackenzie & Joe Argentino, Memorial University of Newfoundland

Commonalities between musical and phonological structure will be illustrated through an examination of reversing operations found in language games and pitch patterns found in serial music. Both music and language constitute complex and abstract systems which are human-specific. A body of literature has developed investigating the similarities of music and language with respect to structural properties of the two systems (e.g. Lerdahl & Jackendoff 1983, Katz & Pesetsky 2011). Our talk contributes to research on the formal similarities between language and music by focusing on phonological language games and serial music.

Phonological studies of language games (e.g. Bagemihl 1989, 1995, Vaux 2011) have provided evidence for theories of phonological representations. Constituents manipulated by language games are abstract, theoretically motivated units such as syllables and timing units. These constituents have also been invoked in work exploring similarities between musical and linguistic structures, particularly in analyses of metrical structure necessary for stress assignment and grouping structures in music (e.g. Fabb & Halle 2011). This talk presents additional evidence for hierarchical constituent structure in music by demonstrating commonalities between manipulation of prosodic constituents in language games and related pitch sequence structures in serial music. Furthermore, we show parallels not only between musical and phonological representations, but also between the language game operations themselves and the manipulation of pitch-class segments in serial compositions; a compositional technique in which pitches are fixed with a specific ordering that is subject to rearrangements.

Language game patterns are compared with pitch collections used in serial music. Pitch collections can be transformed so that the order of pitches changes but the content of the pitch collections remains the same. Manipulations of pitch collections parallel patterns found in language games crosslinguistically. We present pitch collection manipulations which mirror attested language game operations including syllabic reversal, transposition, and interchange.

The phonological example below shows a transposing game in Tagalog (Bagemihl 1989). When the Tagalog form /kamatis/, ‘tomato’, undergoes the game, the final syllable is moved to the beginning of the word, giving /tiskama/. The musical example shows a pitch transformation from Schoenberg’s *Modern Psalm* Op. 50C in which the final two pitches (i.e., dyad) are moved to the initial position. This transformation operates over an abstractly defined pitch sequence in which individual pitches may be repeated without altering the sequential structure.

(1)

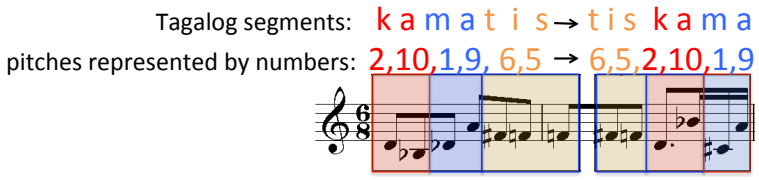

The transformations in both examples involve moving material from the end of the structure to the beginning. Both transformations involve manipulation of an abstractly defined subconstituent of the original structure. In the game, the final syllable is moved to the initial position of the word but the order of segments within the syllable remains constant. In the musical transformation, the final dyad is moved to the initial position while the order of pitches within the dyad remains unaltered. Parallels of this type provide evidence for commonalities between the representation and computation of musical and linguistic structures as well as between phonological patterns and patterning in other areas of cognition more generally (e.g. Moreton & Pertsova 2012).

What we didn't think we could induce from low-pass filtered speech

Patrick Mross & Dinah Baer-Henney // University of Düsseldorf
{Patrick.Mross; Dinah.Baer-Henney}@uni-duesseldorf.de

In a number of papers Berent and several colleagues argued that speakers have knowledge of patterns, even if they have not heard these patterns in the surrounding language (Berent et al., 2007; Berent & Lennertz, 2007). English speakers know that in word-initial clusters rises are less marked than plateaus which in turn are less marked than falls. The authors argue that this knowledge cannot be acquired, since the evidence of markedness is not explicitly available in the input. Moreover, Gómez et al. (2014) suggest that even newborns are aware of markedness of onset clusters. They argue that it cannot be maintained that these children have learned about onset phonotactics on the basis of the input.

It must be noted, though, that after seven months of gestation the auditory system is working and unborn children already hear filtered speech, which may contain all the cues children need to start learning about the phonotactics of syllables. The speech signal available to the child is least informative with respect to sonority profiles. Research that simulates this situation usually uses manipulated speech which has been low-pass filtered at 400 Hz (Mehler et al., 1988; Nazzi et al., 1998). Thus the question remains whether a classification of sonority profiles is possible on the basis of filtered speech. Does the unborn child already receive enough information to learn about sonority profiles? Our study aims to provide a first insight into this issue by testing how much information is still available in low-pass-filtered speech.

In an ABX task we presented 20 German adults with the same types of stimuli used in Berent et al. (2007). Most onset clusters in German have a sonority rise. In the stimuli there were items with a sonority rise, a plateau or a fall. A native speaker of Russian recorded the items and we low-pass-filtered them at 400 Hz. In the experiment we compared the manipulated stimulus with two non-manipulated forms: the target as well as another candidate which differed in the type of sonority profile. We asked participants to identify the filtered stimulus as one of the non-filtered ones. Participants did surprisingly well in this identification task and correctly identified initial falls at a rate of 78.8%, rises at a rate of 76.1% and plateaus at a rate of 71.3%, see fig.1. All sonority profiles are thus identified well above chance level.

We show that there is more information available than assumed before in the signal available to the unborn child. We believe that spectral information, which is informative for categorizing sonority profiles, is still available in low-pass-filtered speech. In short, knowledge about sonority profiles available right after birth may still be a consequence of prenatal learning and does not necessarily need to be assumed to be innate.

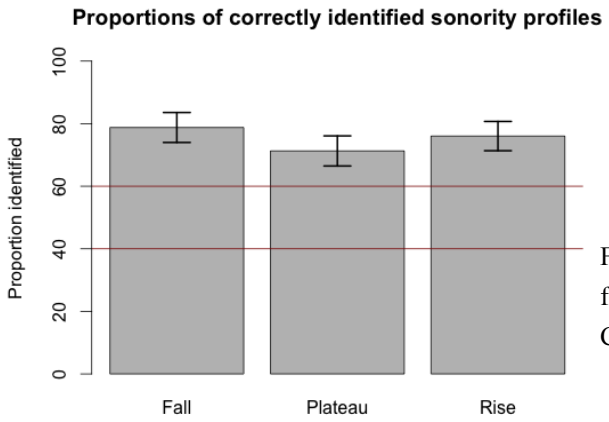

Figure 1: Proportion of correctly identified sonority profiles by cluster types. Chance level marked by red lines.

The more the merrier: A-sharing in the domain

Filiz Mutlu, Boğaziçi University, filiz.mutlu@boun.edu.tr

Aim: A unified analysis for vowel-dependent strident affrication, mid-vowel lowering before rhotics and retroflexation next to rhotics is presented. I claim that if a domain (two objects consistently restricting the shape of one another) has more than one instance of the element **A**, their interaction yields a new object (affrication, etc).

Background: My framework is Government Phonology (Kaye et al. 1985, 1990). Coronals contain the element **A** (and possibly other elements), rhotics only **A** (Broadbent, 1991). **A** has been reinterpreted as structure (Pöchtrager, 2015). An element occurs in a phonological expression only once, but structures can recur (up to a maximum of four layers).

P**roblem I**: Quebec French has [ts, dz] before [i, ɪ, y, ʏ] and [t, d] elsewhere (Kaye, 1989). Japanese has [tʃ, dʒ] before [i], [ts] before [ɯ] and [t, d] elsewhere (Yoshida, 1996).

Analysis: A coronal obstruent stop+nucleus is a domain. If the stop head is a complex structure **AA** (strident), the nucleus must not have **A**. If the nucleus has **A**, the stop head must be only **A** (plain alveolar). There are exactly two instances of **A**. This accounts for the complementary distribution of strident stops and **A**-nuclei: [tsɯ] but [to]. The element **I** is required for **A**-interaction in Quebec French: [tsy] but [tu]. Japanese is not thus restricted.

Problem II: Turkish has [ɛ] not [e] before rhotics in closed syllables: [ɛr] 'man' but [eri] 'man.ACC'. Assuming schwa to be invisible to syllable structure (Wiese, 2000), we can say Swiss German short mid vowels before rhotics are always open in closed syllables (Sutter, pc.).

xærblə 'to puke' *xærb* 'basket.PL' *ʃɔrə* 'to paw' *xɔrn* 'grain' *xɛrn* 'pip'

Analysis: Two **A**'s strictly adjacent in the same domain cannot behave independently but must interact. A closed syllable is a domain. Remember rhotics are **A** alone, other coronals contain more elements. If no other element intervenes, **A** is adjacent to the nucleus. **A** gets copied iff adjacent to **A**. Similar to Qu. French, the element **I** is required for **A**-interaction in Turkish: [ɛr] but not [ɔr].

Problem III: Norwegian [ɾ]+[t, d, n, s] ⟶ [ʈ, ɖ, ɳ, ʂ] in the same prosodic domain (Kristoffersen, 2010). [s] retroflexation is optional, the rest obligatory (Johnsen, 2012).

Analysis: A has two available points. **A** merging with **A** creates variation on the theme of coronal depending on where merge happens. Illustration of different kinds of coronal heads:

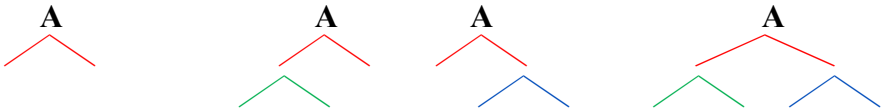

plain alveolar strident retroflex strident retroflex

Onset **A** and rhotic **A** must merge iff both are simplex. **AA** [s] merges with **A** optionally.

Unified analysis: Affrication, retroflexation and mid-vowel lowering are the result of **A** interaction. They differ by the number of **A**'s in the domain.

onset **A** ⟷ nuclear **A** : strident affrication: Domain must contain exactly two **A**'s.

rhotic **A** ⟷ nuclear **A** : mid-vowel lowering: Domain must contain more than two **A**'s.

rhotic **A** ⟷ onset **A** : retroflexation: Domain must contain at least two **A**'s.

Affrication and lowering are dependent on the element **I** in Quebec French and Turkish. The affrication domain is onset+nucleus, the lowering domain nucleus+onset. Retroflexation is not dependent on **I** or onset/nucleus ordering: In Scottish English "initial /tr/ and /dr/ are often also post-alveolar affricates, but *rhoticised* ones" (Scobbie at al., 2006 emphasis mine). At this point it remains unclear why and how **I** interacts with **A** and syllable structure.

Nasal cluster dissimilation in Ngarinyman

Stephen Nichols, University of Manchester, stephen.nichols@manchester.ac.uk

Ngarinyman is an Ngumbin (Pama-Nyungan) language spoken in Australia's Northern Territory which exhibits a phenomenon known as nasal cluster dissimilation (NCD). This is very similar to NCD found in the closely related Gurindji language (McConvell 1988, 1993). Indeed, NCD is the 'most pervasive type of dissimilation' found in Australian languages (Dixon 2002:265). As shown in (1), in Ngarinyman an underlying nasal-plosive cluster is reduced to a simple plosive when preceded by another nasal-plosive cluster (Jones 1994:12).

(1) a. /jaɖi-ŋga/ → [jaɖi-ŋga] 'shade-LOC' (Jones 1994:12)
b. /gaɳdi-ŋga/ → [gaɳdi-ga] 'tree-LOC' (Jones 1994:20)

This is reminiscent of Meinhof's Law and especially the Kwanyama Law found in certain Bantu languages (Meeusen 1963, Johnson 1979, Herbert 1986, Kim 1999, inter alia), as illustrated in (2) and (3) with data from Herbert (1976) (reproduced from Stanton (2015a)).

(2) *Meinhof's Law in Kikuyu*
a. ro-reme → neme 'language(s)'
b. ro-ɣɛɛndɔ → ŋɛɛndɔ 'journey(s)'

(3) *The Kwanyma Law in Kwanyama*
a. oŋgandu → oŋgadu
b. ombambi → ombabi

Here I set out a constraint-based analysis of NCD in Ngarinyman which treats clusters as incidental. This centres around the high-ranked anti-oscillation constraint *[NCNC]$_{\omega}$ which penalises the sequence [+nas][-nas][+nas][-nas] at the word level, such as in (4).

(4) * g a ɳ d i - ŋ g a
[-nas] [+nas] [-nas] [+nas] [-nas]

As only consonants specified as [-cont] are overtly specified for [±nas], blocking is imposed only by intervening [-cont] but not [+cont] consonants. NCD can therefore superficially occur at long distance across certain segments but all instances of NCD are in actuality local. This conclusion is also reached for Gurindji by Stanton (2015b). The resolution that wins out is explained by the relative ranking of the remaining constraints:

1. *[NCNC]$_{\omega}$: penalise the sequence [+nas][-nas][+nas][-nas] at the word level
2. IDENT(root): penalise changes made to the root
3. IDENT(onset): penalise changes made to onsets
4. *|CC|: penalise morpheme-internal geminates
5. MAX: penalise deletion

An example of their implementation is provided in (5) below.

(5) /gaɳdi-ŋga/	*[NCNC]$_{\omega}$	IDENT(root)	IDENT(onset)	*|CC|	MAX
[gaddi-ŋga]		*!		*	
[gadi-ŋga]		*!			*
☞ [gaɳdi-ga]					*
[gaɳdi-gga]				*!	
[gaɳdi-ŋa]			*!		*
[gaɳdi-ŋga]	*!				
[gaɳdi-ŋŋa]	*!		*	*	
[gaɲi-ŋga]		*!	*		*
[gaɲɲi-ŋga]		*!	*	*	

Restricting NCD to the level of the phonological word also explains the variability of the auxiliary /ba/ which may behave either as a clitic or an independent word (Jones 1994). If *[NCNC]$_{\omega}$ is expanded to include the avoidance of [-nas][+nas][-nas][+nas], the analysis may also explain apparent nasal-plosive alternations found elsewhere, outside of cluster contexts.

In this paper, I provide a coherent constraint-based explanation for the behaviour of NCD in Ngarinyman which does not require giving any special treatment to nasal-plosive clusters.

The inherent vowel prime in Fijian

Hitomi Onuma
Iwate Medical University

Kuniya Nasukawa
Tohoku Gakuin University

Koizumi Masatoshi
Tohoku University

Fijian employs *i* as its default epenthetic vowel (Schütz 1978, 1985; Kenstowicz 2007; Kumagai 2014). To confirm this, we investigate what kind of vowel epenthesis is observed when Fijian creates a novel loanword by borrowing from English (rather than established loanwords that are listed in dictionaries, as it is not clear from the data based on established loanwords whether a word has been nativised in Fijian – and if so, when – or whether a word has been borrowed directly from English or indirectly via another language such as French.). The patterns examined are all based on data that we collected from four native Fijian informants in Fiji during August 2014. The four informants were given approximately 400 English words and asked to respond by reproducing these words in a nativised form.

Focusing in particular on epenthetic vowels between consonants and after word-final consonants, the results obtained from the four informants are shown below. (the data does not include any vowels which were copied from source words in the donor language, English.)

Epenthetic Vs	a. Informant A		b. Informant B		c. Informant C		d. Informant D	
	i	49.7%	*i*	57.2%	*i*	70.1%	*i*	69.4%
	e	35.4%	*e*	28.4%	*e*	14.6%	*e*	22.5%
	a	8.0%	*a*	2.4%	*a*	9.4%	*a*	5.0%
	o	2.2%	*o*	4.6%	*o*	1.4%	*o*	0.5%
	u	4.5%	*u*	7.1%	*u*	4.2%	*u*	2.5%

According to our investigation, among the five vowels in the system, the vowels *a*, *o*, *u* are the ones that appear to have been influenced by the quality of adjacent consonants (e.g. in the word 'approve' the informants gave the nativised pronunciation *apuru:βe*, where we assume the first *u* comes as a result of being influenced by the place feature of the preceding consonant *p*). However, limiting the present discussion to epenthetic vowels that are context-free, the data shows a strong preference for the front vowels *i* and *e* as epenthetic vowels in Fijian. Between these vowels, *i* is maximally unmarked, underspecified, and is phonetically the shortest and the least peripheral in the vowel space (Kenstowicz 2003, Uffmann 2006).

This may explain why they most highly select *i*. However this kind of explanation cannot account for why *e* is secondly selected. Given this, an argument should be phonological rather than phonetic-oriented description. A possible analysis may be provided by referring to phonological primes. Within the context of Element Theory, it can be said that its behaviour is attributed to the status of the element |I| in Fijian: |I|, as opposed to |A| or |U|, is the inherent feature serving as the base for constructing phonological structure in Fijian. The phonetic manifestation of a sole |I| (with lighter timbre) is the high front vowel *i*. Then, this tells us why *e* is the second most frequent epenthetic vowel in the language: *e* contains |I| in its internal structure (Backley 2011).

Typologically speaking, the same inherent role is played by |U| in Japanese and |A| in English: the inherent vowel |U| phonetically manifests itself as *ɯ* in Japanese, while inherent |A| is realised as *ə* in English. These two vowels function as default epenthetic vowels in loanwords in the respective languages. Identifying the inherent vowel feature helps us analyse vowel-driven phonological phenomena in a given language.

The existence of morphologically conditioned phonological shifts has engendered a number of proposals to formalize phonological scales, e.g. Kirchner 1996, Gnanadesikan 1997, Mortensen 2008, Pycha 2008, and others. Such proposals allow for incremental shifting in a single direction (chain-shifts) in synchronic grammar. Mortensen draws attention to the problem of neutralization, in which a chain shift collapses contrasts at the far end of the scale. One well-known response to this is the famous circle-shift phenomenon in languages like Xiamen. Another, not previously documented, occurs in Guébie, an Eastern Kru language spoken in southwest Côte d'Ivoire, where the contrast between perfective and imperfective aspect is marked on the verb by a scalar tone shift (1).

Tone is marked with numbers 1-4 where 4 is high, and verbs are bold.

(1) **Morphological tone changes: perfective vs. imperfective**

	Perfective			**Imperfective**	
a.	ɔ^{3} **gbala**$^{4.2}$ si^{2}	'He climbed trees.'	b.	ɔ^{3} **gbala**$^{3.2}$ si^{3}	'He climbs trees.'
c.	ɔ^{3} **li**^{3} ɟa^{31}	'He ate coconuts.'	d.	ɔ^{3} **li**^{2}	'He eats coconuts.'
e.	ɟaci$^{2.31}$ **pa**^{31} gɔlɔ$^{3.3}$	'Jachi flipped the boat'	f.	ɟaci$^{2.31}$ **pa**^{21} gɔlɔ$^{3.3}$	'Jachi flips boats.'

I consider the perfective tone of the verb, (1a,c,e), to be *basic*. Perfective tone matches the tone of the verb in SAuxOV sentences, where the auxiliary and not the verb encodes aspect. Tone on the first syllable of imperfective verbs in SVO sentences surfaces one step lower on the four-tone scale than the corresponding perfective verb, (1b,d,f). If the first syllable contains a contour tone, only the first level of the contour is affected, (1f).

(2) **Perfective to imperfective tone changes**
$4 \gg 3 \gg 2 \gg 1$

Some perfectives have a low tone (scalar value 1). Since imperfective aspect is realized by tone lowering, we might expect either that tone 1 perfectives will surface as superlow (scalar 0) in the imperfective, or that they will be tonally identical across the aspects. Neither of these predictions holds. Instead, contrast between perfective and imperfective low-toned verbs is maintained by raising the final tone of the preceding word (the subject) by one step, even if this results in a superhigh tone, tone 5 (3b).

(3) **Imperfective is derived from perfective**

a. e^{4} **g^{w}rile**$^{1.1}$ ɔ^{2} 'I cursed him' b. **e**^{5} **g^{w}rile**$^{1.1}$ ɔ^{2} 'I curse him'

Thus, the aspect-marking scalar tone changes in Guébie not only affect the verb, but also the tone on the subject, (3b), just in the case that not doing so would lead to ambiguity between perfective and imperfective forms. I claim that these facts can be most insightfully captured by an Agreement-by-Correspondence analysis (Hansson 2001, Rose & Walker 2004), where the tones of the subject and verb are in correspondence, and the imperfective context triggers dissimilation (Bennett 2013) requiring the subject and verb to be one step further apart tonally in the output than the input. A preference for tone lowering rather than raising, and faithfulness to nouns over verbs derives the typical verb tone lowering in imperfective contexts. However, when this is impossible, the preference to maintain contrast is acquired by raising the subject tone. These Guébie facts illustrate the crucial role that paradigmatic contrast maintenance can play in phonology (Lubowicz 2003, McCarthy 2005).

Neural Correlates of Voicing Specification in English and Arabic Fricatives

Kevin Schluter[1], Stephen-Politzer-Ahles[2,1], Meera Al Kaabi[3,1], Diogo Almeida[1]
[1]New York University Abu Dhabi, [2]University of Oxford, [3]United Arab Emirates University

Recent investigations have used neurophysiological evidence – the mismatch negativity (MMN) response – to test for the underspecification of phonological features (Eulitz & Lahiri 2004, Lahiri & Reetz 2010, among others). Unfortunately, investigation of the well-studied [CORONAL] feature is problematic because it is hypothesized to be a universal feature (Avery & Rice 1989, among others) and its prior effects may be related to simply universal acoustic or neuronal properties. A difference between distinct language populations with different featural specifications for the same or similar sounds would provide strong evidence for the phonological nature of this response. We utilize the laryngeal realist hypothesis (Iverson & Salmons 1995, Honeybone 2005) to test English and Arabic-speaking populations with fricatives (s, z; θ, ð) that may have distinct phonological specifications in each language.

Using a passive oddball paradigm to evoke the MMN response (680 standards, 120 deviants) with a 32-channel EEG, we test two contrasts (s & z, θ & ð) using tokens produced by two native speakers (English, Arabic), with two subject populations (English monolinguals, Arabic-dominant bilinguals). Multiple tokens of each type were used, and participants watched a video with subtitles in their native language. We hypothesize that if Arabic and English are specified in opposite ways ([VOICE] for Arabic vs [SPREAD GLOTTIS] for English) that we should see opposite asymmetries for each subject population.

A visual inspection of preliminary means and 95% confidence intervals (18 Arabic and 10 English speakers) suggests that both English and Arabic speakers show a larger MMN for the /s/→[z] condition than the /z/→[s] condition (see figure below), though the magnitude of the effect may differ. This is the pattern predicted for [SPREAD GLOTTIS] specification. The /θ/ and /ð/ tokens are not yet conclusive, but trend in the same direction (not pictured). If this pattern holds, one of two conclusions is likely. The first interpretation of the data is that Arabic and English fricatives have the same phonological specification, namely [SPREAD GLOTTIS] is the specified feature. The second interpretation is that this experiment was unable to access a phonological representation for Arabic, perhaps because the Arabic speaking subjects are – by necessity – bilingual in English.

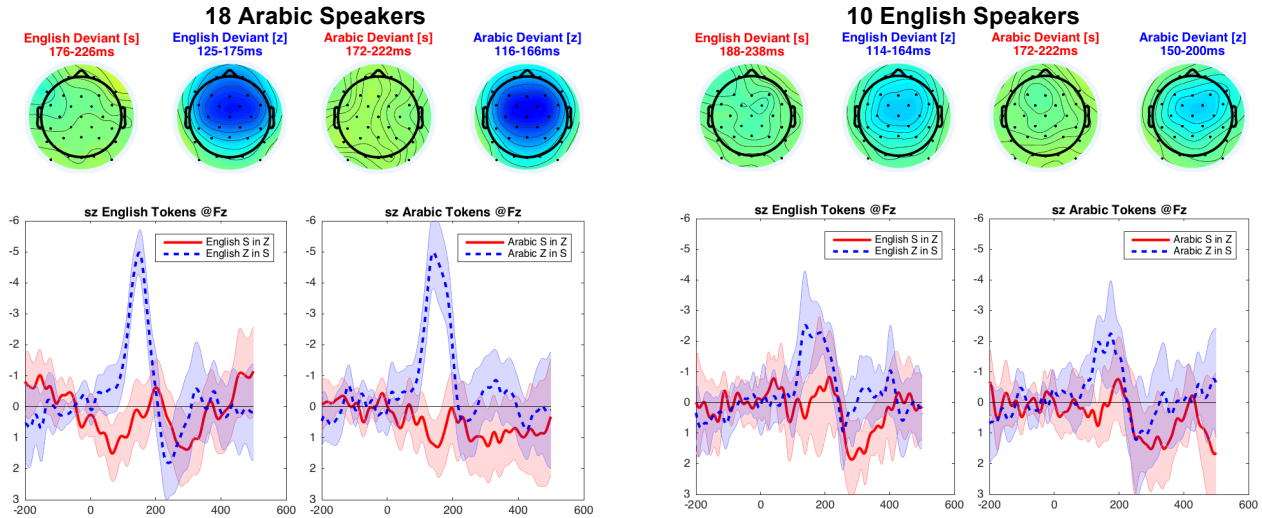

Waveforms at Fz and topoplots for Arabic (Left) and English (Right) speakers for both Arabic and English tokens of [s] and [z]. In all cases the blue line represents deviant [z] and the red line deviant [s]. A large effect for [s] is in line with the predictions of [SPREAD GLOTTIS] being the specified feature.

Trigger deletion in Gurindji

Juliet Stanton, MIT – juliets@mit.edu

It is generally accepted that unbounded spreading is myopic (Wilson 2006, McCarthy 2011, though cf. Walker 2014); i.e., spreading processes cannot look ahead, in the following sense: given a regressive spreading process for some feature [*F*], and a domain [*w x y z*], the decision to spread [*F*] from *z* to *x* does not take into account whether [*F*] will succeed in spreading to *w*, the domain's edge. This paper argues, however, that a non-myopic spreading process is attested in Gurindji (Pama-Nyungan, McConvell 1988). When full application of an unbounded nasal harmony process would violate a local phonotactic, the nasal trigger deletes, blocking spreading from occurring altogether. The existence of the Gurindji pattern has implications for theories of of harmony and the structure of the phonological grammar more generally, as the only theories capable of accounting for it are those theories which allow global evaluation of surface candidates (e.g. parallel Optimality Theory; Prince & Smolensky 2004).

Data. In Gurindji, sequences of nasal clusters ($NC_1…NC_2$) are dispreferred. The observed repair differs according to NC_2's composition: if it is homorganic, N_2 deletes (*kanyju+mpal* > *kanyju-pal* 'across below'); if heterorganic, N_2 denasalizes (*nyampa-n-pula nya-nya* > *nyampa-t-pula nya-nya* 'what did you two see?'). The application of these processes is unbounded, but constrained by the material intervening between the two NCs. If the intervening material contains only continuants (e.g. [w], [l]), N_2 modification is obligatory. If it also contains a non-continuant (e.g. [t], [p]), however, N_2 modification is blocked. Illustrative examples follow. (A regular process leniting intervocalic /k/ to [w], observable below, is not discussed here. Not pictured below is the fact that onset nasals also block harmony; this point will be discussed.)

Intervener	N_2 modified?	Example
…l…	Yes	*kankula+mpa* > *kankula-pa* 'on the high ground'
…w…	Yes	*yangki+kumpalng* > *yangki-**w**upalng* 'to avoid asking'
…t…	No	*nampiyi**t**a-wunyja* '(animal) lacking a female'
…p…	No	*ngu-ngantipa-ngkulu nya-nya* 'they saw us'

Proposal. I argue that the data should be interpreted as follows. Assume that in the normal case, Gurindji has regressive [+nasal] harmony, triggered by coda nasal consonants. [+nasal] can spread through all continuant segments. Non-continuants block spreading (so orthographic *nampiyita-wunyja* should be transcribed as [nãmpijitã-wũnɪa]). [+nasal] spreading does not occur, however, when full application would result in the nasalization of a post-NC vowel (e.g. *[kãnkũl̃ã-mpa]; see Stanton to appear on a perceptual motivation for this restriction). The response to this restriction is not to spread [+nasal] partway (*[kankul̃ã-mpa]), but to prevent spreading from occurring altogether by destroying the [+nasal] trigger ([kankula-pa]). Several alternative interpretations of the data (e.g. [-nasal] spreads progressively from NC_1, causing N_2 modification; N_2 modification occurs to satisfy a constraint banning multiple NCs within a word) will be considered and rejected, on both language-internal and typological grounds.

Nasal spreading in Gurindji is thus non-myopic: given a domain [*w x y z*], whether or not [+nasal] spreads from *z* to *y* crucially depends on whether or not it succeeds in spreading all the way to *w*. I show that this pattern receives a natural analysis under theories in which well-formedness is globally evaluated (e.g. parallel Optimality Theory, Prince & Smolensky 2004), but that it cannot be analyzed in models that restrict the scope of the evaluation to adjacent segments, thereby excluding the possibility of non-myopic processes (e.g. local iterative spreading; the targeted constraints of Wilson 2006; Harmonic Serialism as in McCarthy 2011).

Outlook. Non-myopia exists beyond Gurindji: harmony processes in Baiyina Oroquen and Mọ̀ba Yorùbá involve trigger-target relations that are demonstrably non-local (Walker 2014), and an across-the-board [a]-raising process in Romanian is initiated only if all resulting [ʌ]s are phonotactically licit (Steriade 2016). This small but growing class of non-myopic patterns suggests that any successful theory of unbounded spreading should predict their existence.

The Accent Locality Hypothesis and parameter dependency: toward a perfect fit

Alexandre Vaxman
University of Connecticut
alexandre.vaxman@uconn.edu

This talk presents important refinements to the Primary Accent First theory ("PAF") proposed in van der Hulst (1996, 1997, 2010, 2012), *a.o*. This theory separates word accent (primary stress) from rhythm (non-primary stress) and assigns those independently. It attempts to capture cross-linguistic variation in accentual patterns in terms of a small number of parameters. While, in many cases, PAF makes correct predictions, it also significantly overgenerates. Thus, as I will show, of the 36 weight-sensitive (WS) systems generated by the PAF parameter system, 19 are unattested.

The goal of this paper is to reduce the parameter space of the PAF grammar, while retaining PAF's correct predictions. To that end, I introduce a particular parameter dependency into the PAF grammar and modify the Extrametricality parameter.

First, since initial extrametricality is cross-linguistically very rare, the Extrametricality parameter (Left/Right) is replaced with the Nonfinality parameter (Yes/No). This rules out initial extrametricality altogether. (Nonfinality (Yes) results in final extrametricality).

Second, I propose (1) for WS systems with nonfinality:

(1) *The Accent Locality Hypothesis*

In a WS system with nonfinality, accent in words with heavy syllables must fall on the heavy syllable closest to the right word edge.

This hypothesis is theoretically desirable because it sets a strong locality restriction on accent location by minimizing the distance between the accented heavy syllable and the right edge of the word (2).

(2) l h l l (h ˈh) <σ>] *l h l l (ˈh h) <σ>]

Traditionally, accent is assigned in PAF by the Select parameter, which chooses the {Left, Right} heavy syllable in the accent domain. However, as (1) implies, Select must be fixed to the "Right" setting in WS systems with nonfinality, which makes the testable, falsifiable prediction (3).

(3) *The combination {Nonfinality (Yes), Select (Left)} is unattested.*

Testing (3) against the data in StressTyp, the largest database of stress patterns of the world's languages (van der Hulst *et al.* 1996) reveals that the prediction (3) is borne out, thus supporting the Accent Locality Hypothesis (1). This empirical result justifies the dependency of Select on Nonfinality.

The changes to the PAF grammar proposed here successfully eliminate all the unattested combinations of parameter settings, while leaving the attested ones unaffected, bringing the revised PAF grammar very close to the level of descriptive adequacy.

Assimilation-Driven Integrity in Clusters

Islam Youssef <islam.youssef@hit.no>
University College of Southeast Norway

This talk addresses the lack of epenthesis in Baghdadi Arabic (BA) within consonant sequences that undergo assimilation vis-à-vis other sequences. Whereas BA final clusters are broken up by epenthesis (1a) (Blanc 1964), a word-final true geminate stays intact (1b) – a fact attributed to *geminate integrity* (Hayes 1986). On the other hand, false geminates arising from morpheme concatenation are known to induce epenthesis (1c) (Majdi & Winston 1993).

(1) a. /ʔibn/ → [ʔibin] 'son' b. /sitt/ → [sitt] /*[sitit] 'lady'
c. /fut-t/ → [futit] 'I passed' (a *t*-final stem attached to first person singular suffix -*t*)

False geminates arising from total assimilation resist epenthesis just like true geminates, and behave as if the morpheme boundary has been obliterated. For example, BA four- or five-consonant sequences are broken up as CCiCC and CiCCiCC, and if the last two consonants constitute a geminate (GG), epenthesis applies regularly. However, when a geminate appears immediately before the last consonant, epenthesis applies before and/or after but never internal to the geminate: CiGGC or CCiGGiC (Rose 2000). As exemplified in (2a), the expected form *CGiGC in which the epenthetic vowel splits the false geminate (after total assimilation of the definite article *l-*) is not attested. Interestingly, heterogeneous clusters resulting from partial assimilation are also immune to epenthesis. In (2b), the final cluster resulting from nasal place assimilation is either assimilated or split up by [i], but never both.

(2) a. /ħisaːb **l-sn**iːn/ → [ħi͡saːb͡ i͡ **ss**i͡**n**iːn] /*[ħisaːb **s**i**sn**iːn] 'counting the years'
b. /ʤa**nb**/ → [ʤa**mb**] / [ʤa**ni**b] /*[ʤa**mi**b] 'beside'

I argue that the lack of epenthesis in these sequences is attributed to assimilation, which results in a doubly linked structure and some discrepancy with the underlying form. The essence of the analysis is that assimilatory feature-linkage provides a consonant sequence with immunity against epenthesis. If this pertains to false geminates formed by across-morpheme-boundary assimilation, it becomes unnecessary to treat outputs as true geminates (where geminate integrity operates) or to assume deletion of morpheme boundaries, as in Guerssel (1978). And since both partial and total assimilation involve action at the level of the individual feature, a unified autosegmental account (Goldsmith 1976) is attainable.

While the opaque interaction of epenthesis and assimilation poses a challenge to a derivational model of phonology, a non-derivational model like Optimality Theory (Prince & Smolensky 1993/2004) allows for an account of these multiple operations within a single constraint ranking. First, the fact that epenthesis generally splits final clusters entails ranking *COMPLEX CODA below DEP-V. Second, the fact that epenthesis does not affect true geminates suggests that NOGEM is also ranked below DEP-V (Baković 2005). Next, to split false geminates of morpheme concatenation, we need a highly ranked positional markedness constraint, NOGEM/MORPH, which disallows geminates across morpheme boundaries. Finally, this latter constraint must be outranked by some assimilation-driving markedness constraint, and by another assimilation-integrity constraint that preserves newly shared features. The general ranking in (3) captures the resistance of all assimilated sequences, both geminate and non-geminate, to epenthesis.

(3) LINK [F_i]/DOMAIN, LINK [F_i]-INTEGRITY >> NOGEM/MORPH >> DEP-V >> DEP [F_i], NOGEM, *COMPLEX CODA

OPACITY HAS ONLY TWO SOURCES: THE LEXICON AND CYCLIC SPELL-OUT

Sławomir Zdziebko
John Paul II Catholic University of Lublin

Available analyses of Polish opaque vocalic alternations rely on three analytical devices: (i) direct reference to morphological information; (ii) diacritic marking of vowels by means of feature [+/-tense]; (iii) extrinsic rule ordering. In this paper I will argue that these three conceptually questionable devices make wrong predictions about the vocalic alternations attested in Polish. Instead, I propose that the opacity is the function of the shape of vocabulary items and the availability of the contextual information, i.e. the lexicon and cyclic spell-out.

Thus the rule of Derived Imperfective Tensing (Gussmann 1980:76) may be reinterpreted as integration of an autosegment which re-writes a functional verbal structure containing the Secondary Imperfective (SI) head as in (1).

(1) {V,SI} [aperture[A]] / ___ v[Class Ø/-*i*-]

A combination of constraints which force the replacement of the underlying vowels makes (1) successfully account for the /ɔ/ /a/ and /ɛ/ /a/ alternations (*wyr<u>o</u>bić - wyr<u>a</u>biać* 'to manufacture - SI' and *zgni<u>e</u>ść - zgni<u>a</u>tać* 'to smash - SI').

Importantly, verbs which realize the SI-head as affix *-iw-* and verbs which are not Secondary Imperfectives but contain the affix *-aj-* do not show the relevant alternations. Traditionally, this has been achieved by the direct reference to the affix *-aj-* and the label Derived Imperfective in the environment of the rule. In the current account no such reference is necessary. Verbs which do not show the alternation realize the root, V and SI-heads by means of the stem exponent. It is, therefore, the size of vocabulary item that the presence of the alternation depends on. Exponent *-aj-* realizes higher functional structure.

Polish shows an alternation of thematic /a/ and /ɛ/ in the masculine-personal gender of verbs in *-ej-*: *siwi<u>a</u>łyśmy - siwi<u>e</u>liśmy* 'we went gray, non-masc-pers. - masc-pers.'. In the traditional account of the alternation (Gussmann 1980:65) tense /ɛ/ undergoes Backing and Lowering (BL) if followed by a [+back] consonant. Hence BL must crucially be ordered after the rule of Palatalization which takes [+back] /ɬ/ into [-back] /l/.

Gussmann's account is falsified by a productive class of resultative adjectives derived from past unaccusative verbs where masculine-personal forms invariably show /a/: *osiwi<u>a</u>łe - osiwi<u>a</u>li* 'gone gray, non-masc-pers. - masc-pers.'. The analysis which derived the correct output in the case of verbs fails as the early ordering of Palatalization should bleed BL.

As in the case of stem-vowel alternations I will argue that the /a/ /ɛ/ alternation is a consequence of the shape of vocabulary items that realize the functional structure in the relevant verbs and resultative adjectives (2).

(2) {V,v[Class *ej*],ViewAsp} /ej/ / __ Inf/Imp/[+masc-pers]/[-past] (a)
/a/ (b)

(2a) realizes the functional structure in the Infinitive, Imperative, in the context of feature [+masculine-personal] and in the non-past forms of verbs (/j/ is not realized before consonants for phonological reasons). /a/ is the elsewhere exponent.

The structure of de-verbal resultative adjectives involves merging of a category-defining Adj-head above the Tense head. Following Embick (2010) I assume that category-defining heads trigger spell-out of the cyclic domains in their complement and the material merged above the cyclic head is not available when the cyclic domain undergoes Vocabulary Insertion. As a result, feature [+masc-pers], which is part of the Agr-node and dominates the Adj-head, is not available when (2a) and (2b) compete. Thus /a/ is inserted and serves as the thematic vowel in resultative adjectives regardless of the gender specification.

Geminates in the phonetics/phonology interface:
an analysis of the release burst in Polish true geminates

Dariusz Zembrzuski[1], Aleksandra Karwacka[1] and Jakub Szewczyk[2]
d.zembrzuski@gmail.com, akkarwacka@gmail.com, jakub.szewczyk@gmail.com
[1]University of Warsaw, [2]Jagiellonian University

The phonetic analysis of geminates has so far concentrated on various acoustic cues to differentiate these segments from singletons (Esposito, A. and di Benedetto, M. G. 1999; Arvaniti, A., and Tserdanelis, G. 2000; Al-Tamimi, J. and Khattab, G. 2011).
Much less research has been dedicated to the topic of release burst, especially in the context of geminate representation. It is believed that, cross-linguistically, 'true' geminates are typically unreleased. On the other hand, 'fake' geminates exhibit optionality with the use of release burst (Wierzchowska, B. 1971; Kozyra, A. 2015). In this context, it is interesting to ask whether released 'true' geminates are an accidental gap.

To this end, a research has been conducted to check for the presence of release burst in the production of 'true' geminates. The language chosen for the study is Polish, where the representation of identical double consonants appears to present a challenge. Specifically, the generative literature labels them as 'fake' geminates (Rubach, J. 1986), establishing their status on the basis of their behaviour with respect to palatalisation rules. Namely, in palatalised forms, it is only the second segment that undergoes a change. If the first segment assumes the same shape as the following one, this change is explained by independent assimilatory processes (Rubach, J. 1984). However, assimilation is said to target only coronal segments. Dorsal consonants must remain unaltered. To give an example, the word *Mekce* ['mɛktsɛ] (sg. fem. loc. of *Mekka* ['mɛkka] 'Mecca' place name) should never become ['mɛtstsɛ], even in rapid speech. If such a form occurred, the only explanation would be that the double consonant is in fact a 'true' geminate.

To answer our research questions, we tested N=90 Polish speakers with a nonword repetition task based on place names, ten of which contained dorsal geminates. The input list also varied on the presence of plosion after the first segment of a double consonant. The speakers were divided into a younger group (aged 18-30; experimental group) and an older group (aged above 50; control group) due to an expected age effect. The choice of the name type was motivated by the Polish declension pattern. The frame sentence *Jestem w ...* 'I am in ...' required declension into the locative, which contains a palatalising ending [ɛ]. In the experiment, we controlled for several variables, such as age, origin, level of education, and the influence of orthography.

All items with identical double consonants were transcribed and subsequently assessed so as to isolate the group of 'true' and 'fake' geminates. These two categories were compared. Moreover, the productions with 'true' stop plosive geminates were categorised by the presence of release burst and subsequently compared in this respect with the inputs.

The preliminary results of the study reveal that the productions in both age groups contain significantly more 'true' than 'fake' geminates (experimental: $F=9.75$ $p=6.6 \times 10^{-9}$, control: $F=5.26$ $p=2.7 \times 10^{-5}$). The results also reveal that inputs without release burst were more faithfully repeated than those containing plosion ($p<0.0001$). Specifically, the likelihood of mapping the lack of plosion from the input to the output reached 92,9%. On the other hand, the presence of release burst in the input was mapped to the output only with the likelihood of 43,1%. Such results seem to indicate that release burst in geminates is perceptible in the population. Moreover, the use of plosion in geminates seems to be governed by universal principles, according to which the unmarked value for a geminate is that without release burst. Interestingly, the results further show that 'true' geminates can appear with plosion after the first segment. This finding puts into question the representation of 'true' geminates with a single long segment on the melodic tier.

Special session

Evidence in phonology

Categorization, Evidence and Phonology

William J. Idsardi (University of Maryland)

In this talk I will consider various kinds of evidence for abstract categories in phonology, especially in the light of Pierrehumbert (2016). I will first review the findings from computational modelling in Dillon, Dunbar and Idsardi (2013) where we found that a mean shift model of the acoustic 'scatter' of phonemes and their allophones fit the data better than a model for the separate phones. That is, allowing the model to believe that allophones of a phoneme have correlated acoustic properties improves the model fit, as well as the fit with phonological theory. Second, I will briefly review the neuro-imaging findings in Lago, Scharinger, Kronrod and Idsardi (2015) where we found cortical 'warping' toward sound category centers in the perception of sounds intermediate between [s] and [ʃ], consistent with a perceptual magnet effect (Kuhl 1991). Last, I will turn to the question of discontinuous (or discontiguous) categories. Such categories are modelled easily with instance-based methods (exemplars) such as k-nearest neighbors, and there are theoretical (Kronrod 2014) and empirical (Avery and Idsardi 1999) reasons to believe that such categories are possible for speech sounds. I will present on-going work (e.g. Heffner and Idsardi 2015) examining the 'time to learn' in artificial language learning experiments for contiguous and discontiguous categories, showing that discontiguous categories are harder to learn for both speech and non-speech acoustic stimuli. The general conclusion of these lines of research is very familiar, however: all kinds of evidence are eventually relevant; and this is just another illustration of scientific consilience or inference to the best explanation (Whewell 1847, Wilson 1998, Lipton 2003).

Statistics and abstraction in phonology

Janet B. Pierrehumbert (University of Oxford)

In classical generative phonology, the goal of a phonological grammar was to characterize all and only the possible words of a language. This Boolean agenda is readily extended to a probabilistic one. Decades ago, theories of allophonic and sociolinguistic variation already used probabilistic rules or constraints to capture the relative likelihoods, and not merely the existence, of different variants. The regularity of the patterns captured by these rules is an indication that people learn statistical patterns and apply them productively when speaking.

The next step in probabilistic phonological theory is to use inferential statistics to draw conclusions about the underlying representations of words and about the cognitive architecture. In this talk, I will go over some examples of cases in which statistical analysis provides arguments for abstractions. These examples are centred on OCP constraints and constraints on syllable contacts. The analyses all depend on the assumption that people are statistical learners. Key ideas include:

1) Systematicity. Although people can learn variable patterns, they have cognitive biases towards more systematic patterns. Abstract analyses can be justified by showing that they yield more crisp, extreme, or well-differentiated patterns than a more superficial analysis would provide for the same data.

2) Lexical occupancy. Unlikely words containing combinations of infrequent phonological parts can sometimes be found. However, the likelihood that a point in phonological hyperspace is occupied by a real word in the lexicon decreases with the

frequencies of the parts. This means that phonological constraints can be learned by observing what combinations are statistically under-represented. People's well-formedness judgments reflect this fact.

3) Robustness. People in a linguistic community have finite vocabularies and they don't all know the same words. They share the resources of the phonological grammar because these can be reconstructed reliably from differing input.

Applying these simple but powerful ideas, inferential statistics support (at least) three levels of representation for language sound structure: a morphophonological level, a phonological level, and a phonetic level. They provide tools for precise characterization of these levels, as well as for making quantitative predictions about corpus statistics and behavioural data.

How phonology and typology are shaped: lessons from fieldwork

Sharon Rose (University of California, San Diego)

Phonologists draw on a wide range of data from different languages collected by fieldworkers. Choices made by fieldworkers determine the scope of inquiry, and phonologists therefore limit themselves to particular types of data. In this talk, I sketch two areas of inquiry where fieldwork can help broaden the notion of evidence in phonology, drawing primarily from African languages: ideophones and balanced vowel harmony surveys.

Phonologists have sporadically analyzed the class of words known as ideophones, mimetics or expressives (e.g Mester & Ito 1989 on Japanese). However, ideophones are often omitted from descriptions of the prosaic phonology of a language, dismissed as exceptions to generalizations, and/or reported to exhibit 'unusual' phonology. Based on my own and others' fieldwork on ideophones, I argue that the phonology of ideophones is not unconstrained, but provides key evidence concerning the phonology of a language. In some cases, ideophones show a slightly different distribution of segmental contrasts than prosaic phonology. In others, they show expressive lengthening, pitch, phonation and other prosodic factors tied to performance. Ideophones challenge us to widen the scope of what constitutes phonological evidence.

The typological study of vowel harmony has provided the base for some important theoretical proposals. For example, labial harmony is typically parasitic on vowel height and favors gestural uniformity (Kaun 1995). In ATR harmony, high [–ATR] vowels are assumed to be marked due to i) antagonistic articulation of tongue raising and tongue root retraction, and ii) their absence from vowel inventories. Constraints such as 'If [+high], then [+ATR]/not [–ATR]' (Archangeli & Pulleyblank 1994) or *HIGH/[–ATR] in theoretical proposals embody this concept. Yet, Casali (1995, 2002) has argued against this, and Casali (2003) has shown that the inventory assumption is based on certain Niger-Congo languages spoken in West Africa. A typological survey of Nilo-Saharan languages and Niger-Congo languages spoken in East Africa reveals a different picture, in which these vowels are common. Expanding on Casali's survey, I show that high vowel ATR contrasts correlate with ATR harmony in the east; the missing (or allophonic) vowels are mid [+ATR] rather than high [–ATR]. Languages that lack high vowel contrasts tend not to have ATR harmony. Evidence from fieldwork on Nilo-Saharan languages and eastern Niger-Congo languages has not been given due attention. I discuss the implications of these results for the analysis of ATR harmony.