Klaas Seinhorst and Sterre Leufkens

Phonology and phonetics in functional discourse grammar: Interfaces, mismatches, and the direction of processing

Abstract: In this paper we discuss the interfaces between phonological and phonetic representations in Functional Discourse Grammar, and the possible mismatches that occur at those interfaces. Firstly, we discuss different definitions of phonological opacity in the literature, and provide examples with these definitions. We argue that mismatches between phonological and phonetic representations can result from competing pressures of articulatory ease and perceptual distinctivity. In order to model these influences and the resulting mismatches adequately, the model should not be organised strictly top-down: we argue that FDG should incorporate bottom-up influence from the phonetics on the phonology. We show that these influences are language-specific, which entails that bottom-up feedback must involve the Grammatical Component. With this modification of the model's architecture, language users' tendency to speak efficiently can be incorporated into the model, explaining a wide array of phenomena such as (synchronic) reduction, the cross-linguistic frequency of phonological alternations, and (diachronic) grammaticalization.

Keywords: functional phonology, phonology–phonetics interfaces, mismatches, opacity, reduction, grammaticalization

1 Introduction

This article focuses on the phonological and phonetic representations in Functional Discourse Grammar (hereafter "FDG"), their interfaces, and the possible mismatches in which they are involved. We argue that the model, as far as pho-

Acknowledgements: We would like to thank all participants in the workshop "Interfaces in Functional Discourse Grammar" (Oviedo 2019) for valuable comments and discussion, and Daniel García Velasco, Lucía Contreras García, Marieke Olthof, Matthias Klumm, and an anonymous reviewer for their feedback on this paper.

Klaas Seinhorst, Amsterdam Center for Language and Communication, seinhorst@uva.nl
Sterre Leufkens, Utrecht Institute of Linguistics OTS, s.c.leufkens@uu.nl

https://doi.org/10.1515/9783110711592-004

nology and phonetics are concerned, should not be organized strictly top-down: it must allow for bottom-up influences from the phonetics on the phonology. Furthermore, we argue that this bottom-up feedback must involve the Grammatical Component, because these influences are language-specific. With this modification of the model's architecture, language users' tendency to speak efficiently can be incorporated into the model, explaining a wide array of phenomena such as (synchronic) reduction, the cross-linguistic frequency of phonological alternations, and (diachronic) grammaticalization.

This paper is structured as follows. In Section 2, we discuss the notions of transparency and opacity in FDG, as well as their possible motivations; in Section 3, we explain how the notion of opacity is commonly used in the phonological literature; Section 4 treats the phonological and phonetic representations that have been proposed for FDG; Section 5 is concerned with the interfaces within and between phonological and phonetic representations, as well as possible mismatches between these representations. In Section 6 we argue that the model, as far as phonology and phonetics are concerned, cannot maintain its top-down organization: phonetic considerations exert a bottom-up effect on phonological representations. The conclusion remains for Section 7.

2 Opacity in FDG, and its motivations

In this section we provide examples of opaque phenomena, first in general, and second in the phonetic and phonological literature. We also discuss the various possible motivations for different types of opacity to arise.

2.1 Examples of opacity and their motivations

Over the last few years, much work in FDG has been devoted to the notions of transparency and opacity: Hengeveld (2011) provides an introduction of the relevant concepts, Contreras-García (2013) compares the way that different linguistic frameworks deal with transparency, and large-scale typological surveys were done by Leufkens (2015) and Hengeveld and Leufkens (2018). FDG's multi-level architecture is well suited for a straightforward definition of transparency as a one-to-one relation between two elements of different levels of representation, or between two elements within a single level of representation; opacity is the absence of such a relation. We choose to follow Hengeveld and Mackenzie's (this volume) definition, who regard mismatches strictly as numerical deviations of a one-to-one relation. This means that correspondences between elements that are not prototypically related (e.g. between a State-of-Affairs and a noun), or correspondences between discrete elements and continuous representations, do not constitute mismatches under their (and our) definition. Mismatches occur at interfaces, that is, mechanisms of the grammar that execute a set of operations. Hengeveld and Mackenzie recognize Conceptualization, Formulation, Encoding, Articulation and Contextualization as interfaces, but in Section 5 we will add interfaces pertaining to mechanisms operating between and within phonological and phonetic sublevels.

We will illustrate opacity by means of a Dutch sentence that contains (at least) three opaque phenomena.

(1) Mijn zus Sophie woon-t in de binnenstad
POSS.1SG sister Sophie live-PRS.3SG in DEF.COMM city_center
van Praag
of Prague
'My sister Sophie lives in the city center of Prague.'

The first phenomenon is apposition, as seen in the noun phrase *mijn zus Sophie*: there are two Referential Subacts at the Interpersonal Level, but these only correspond to a single Individual at the Representational Level. Such a many-to-one relation is opaque.

The second opaque phenomenon is clausal agreement, which happens within the Morphosyntactic Level. The noun phrase *mijn zus Sophie* is the subject of example (1); in addition, the subject is marked by an agreement suffix –*t* on the verb. Such clausal agreement is opaque, because the same referent is expressed morphosyntactically twice.

A third opaque property is grammatical gender, as seen in *de binnenstad*. Dutch has two nominal genders: common and neuter. The word *binnenstad* has common gender, even though it does not possess any semantic properties that motivate why it should have common or neuter gender. In FDG, this means that a specification at the Morphosyntactic Level has no counterpart at the Representational Level. Such a none-to-one relation is opaque. A language would have transparent syntactic gender if it encoded natural gender distinctions morphosyntactically, or if it did not mark gender at all, like English.

Both in FDG studies on transparency and in other literature, some opaque phenomena have been argued to be motivated by communicative advantages (e.g. Dahl 2004; Barbiers 2008; Trudgill 2009; Leufkens 2015, to appear). This especially holds for phenomena that involve some form of redundancy,[1] i.e. they supply a piece of information multiple times. For instance, in example (1), there are two items that signal the subject of the clause, and while such a structure is not transparent, it does provide the listener with an extra cue to identify the subject correctly. This increases the robustness of the transmission of information, likely increasing the probability of communicative success. Additionally, redundancy has been argued to facilitate processing (e.g. Coles-White 2004; Nichols 2009), increase saliency (Petré 2019), and increase learnability of the redundantly marked feature (e.g. Audring 2014).

However, other forms of opacity clearly lead to a decrease in learnability. For example, grammatical gender is notoriously difficult for language learners (De Houwer and Gillis 1998 and Blom et al. 2008 for Dutch; Van der Velde 2004 for Dutch and French; White et al. 2004 for Spanish): because gender is not predictable in these languages, learners will need to memorize the gender of each individual noun. In a grammatical judgment task of a semi-artificial language with determiner–noun agreement, Ćurčić (2018: 30) found that learners scored correctly more often on noun phrases in which the gender of the noun was motivated biologically than on items where it was not. The same holds for instances of irregularity in verbs, such as the vowel alternations that English strong verbs undergo when inflected for past tense: memorising which verbs undergo which alternation requires an extra effort that is disadvantageous to language users and learners. This type of opaque phenomena emerges when pragmatically or semantically motivated rules grammaticalize over time into purely morphosyntactic rules or features. As such, they have been referred to in the literature as "historical junk" (Lass 1997).

A possible strategy of language users to eliminate opacity is regularization, the elimination of exceptions in favour of regular, predictable structures: this strategy has been attested in the laboratory (a.o. Hudson Kam and Newport 2005; Smith and Wonnacott 2010; Seinhorst 2017). Regularization occurs in natural language too, a classic example being strong verbs that become weak diachronically (cf. Lieberman et al. 2007 for a corpus study of English). The likelihood and speed of this process seems to depend on social properties of the language community: loss of opacity proceeds more quickly in a community with a large L2 learner pro-

---

1 Redundancy can be viewed as a subtype of degeneracy. The latter involves structurally different elements that fulfill the same function, such as the expression of past tense by means of ablaut (speak > spoke) or by a suffix (talk > talked) in English, or the multiple expressions of argument information in the case of argument–verb agreement (Van de Velde 2014). In our interpretation of redundancy, the term only applies to situations in which the structurally different elements occur within the same phrase or clause. Hence, agreement marking is a case of both degeneracy and of redundancy, while past tense inflection in English is a case of degeneracy but not of redundancy.

portion and in situations of language contact, while phenomena like grammatical gender and irregular inflection are more likely to be retained in relatively isolated communities with a large proportion of L1 speakers (e.g. Kusters 2003; Lupyan and Dale 2010; Trudgill 2011). In such languages, opaque features may persist because they do not seem to reduce learnability, despite the absence of a clear communicative or perceptual motivation. It should be noted that languages may also exhibit deregularization, for instance when weak verbs become strong: an example would be the verb *make*, which used to be weak with past tense *maked*. However, we are not aware of any sources that directly compare the effect sizes of both phenomena (regularization and deregularization) with an appropriate statistical analysis.

2.2 Opacity in phonology and phonetics, and its motivation

The examples of opacity in the previous subsection pertain to the Interpersonal, Representational and Morphosyntactic Levels in FDG, but mismatches may occur in the phonology and phonetics as well, for instance at the interface of Phonological Encoding. A phonological surface transcription of example (1), repeated here as (1′) for convenience, would look as (2), with periods indicating syllable boundaries:

(1′)	Mijn	zus	Sophie	woon-t	in	de	binnenstad	van
	POSS.1SG	sister	Sophie	live-PRS.3SG	in	DEF.COMM	city_center	of
	Praag							
	Prague							
	‘My sister Sophie lives in the city center of Prague.’							

(2) /mɛin.zʏs̲o.fi.ʋoːnt.ʔɪn.də.bɪ.nə.stɑt.fɑm.praːχ/

The different nature of the representations (orthographic in (1′), phonological in (2)) makes it somewhat difficult to compare them at first glance, but some differences can be seen, of which we will discuss two here. Firstly, of the two consecutive *s*-es in *zus Sophie*, only a single /s/ remains: this is an example of degemination, a process in which two successive identical consonants are reduced to a singleton segment. The remaining /s/ is ambisyllabic, which we indicated here by underlining it; this means that it is simultaneously the coda of one syllable and the onset of the following syllable. This happens because the Maximum Onset Principle (Kahn 1976; Selkirk 1981) requires segments to be assigned to onsets whenever possible, but since Dutch syllables cannot end in lax vowels such as /ʏ/, the same segment needs to function as the coda of that syllable as well.

Secondly, the final *n* of *van* is realized as a labial nasal /m/, instead of a coronal /n/, because it has assimilated to the place of articulation of the following consonant. Both these phenomena suggest opaque mappings between some sort of lexical representations and their realisations; we will further discuss these representations in Section 3.

The existence of such opaque mappings cannot be motivated only by the same motivations we mentioned above, because many of these phonological processes do not seem to yield any advantages in terms of processing, learnability, or robustness of transmission, nor can they be seen as "historical junk"; there must be another explanation. In functionalist approaches to phonology and phonetics (a.o. Passy 1890; Martinet 1960; Boersma 1998), two forces are assumed to be at play: a pressure towards perceptual clarity, and a pressure towards articulatory ease.[2] This entails that language users prefer unambiguous auditory cues in order to aid successful communication, while at the same time speakers try to expend as little gestural effort as necessary to convey a phonological contrast. These forces counteract each other, as careful speech is typically more effortful than sloppy speech: speakers aim to strike an optimal balance between the two factors, and try to be as efficient as possible in making themselves understood. In addition, we should acknowledge that regressive assimilations are arguably advantageous to the listener, because they anticipate upcoming content and thereby facilitate word recognition: in our example above, the labial place of the nasal signals the presence of a following labial consonant.

The tendency towards perceptual clarity will try to prevent opaque mappings, so such mappings are more likely due to considerations of articulatory effort. The interaction of these pressures is not only situation-specific, but also language-specific: for instance, coronal nasals undergo place assimilation in English, but not in Limburgish (cf. Section 4). Considerable cross-linguistic variation is also found in, for instance, phonotactic restrictions: whereas complex syllable onsets are illicit in many languages, probably because they compromise articulatory ease as well as perceptual distinctivity, Georgian allows for at least six segments in this position.

2 In this paper, we take a dynamic approach to the FDG framework by considering it as a model that reflects the process of the language user. This can be contrasted with a view of the model as primarily describing grammar, that is, reflecting a static version of the language system. For example, Hengeveld and Mackenzie (2008: 2) state that FDG is not a model of the speaker, but "a theory about grammar, but one that tries to reflect psycholinguistic evidence in its basic architecture." Without taking position as to whether FDG should model speakers or grammars, our aim in this paper is to incorporate evidence about the interplay between functional forces in phonology and phonetics, as it exists in the individual speaker–listener, into the FDG model.

It is often assumed that the tendency to be clear co-determines the grammatical choices of a speaker, either within a certain utterance or within the entire sound system, possibly taking the listener's perception process into account (Martinet 1960; Kirchner 1998/2001; Padgett 2003; Hendriks and De Hoop 2001 for semantics). Boersma and Hamann (2008) argue for a non-teleological alternative: language users learn in perception which auditory cues are least ambiguous, and reuse this same knowledge in production.

The interaction of the tendencies towards perceptual clarity and articulatory ease manifests itself in various domains, both at the level of the individual speaker and at the level of the linguistic system. For instance, speakers tend to reduce repetitions of words by, for instance, centralizing vowels and/or deleting segmental content (Koopmans-van Beinum 1980; Ernestus 2000; Johnson 2004); listeners are not able to recognize reduced forms outside of context, but if context is provided they do identify such forms correctly (Kemps et al. 2004). In the structure of sound systems, the maintenance of auditory contrast plays a central role (Liljencrants and Lindblom 1972; Ten Bosch 1991), but only to the extent that sufficient contrast is ensured. Because both pressures exert effects in various domains of the linguistic system, we argue that they have to be integrated in the grammar. We return to this matter in Section 4.

3 The term 'opacity' in the phonological literature

Phonologists would not normally refer to the mismatches discussed in the previous section with the term "opacity". This notion has been discussed extensively in the phonological literature (a.o. Kiparsky 1973; Kenstowicz and Kisseberth 1979), but its definition is more restricted than in FDG-based research (as also signalled by Leufkens 2015: 21–22). In order to understand the difference between what we will call "FDG opacity", that is, the use defined in Section 2.1, and "phonological opacity", that is, the use defined in this section, we need to know a bit more about phonological theory.

Generative models of phonology traditionally assume two levels of representation: an underlying form (UF) and a surface form (SF) (Chomsky and Halle 1968; Prince and Smolensky 1993/2004). The UF is structured in terms of phonemes (i.e. categories that distinguish between different meanings), morphemes, and morphophonemic words; it is the underlying form where the phonological structure of the morphosyntactic representation is retrieved from the lexicon. The SF is structured in terms of prosodic units such as syllables, phonological feet, and intonational phrases; the morphophonemic boundaries that are still present at UF have been erased at SF. The SF is subject to phonotactic restrictions, and in order to ensure that these restrictions are met, repairs to the UF may be required. For instance, the Dutch UF |ɦɔnd+ən| 'dogs' surfaces as /ɦɔndən/, but the singular |ɦɔnd| surfaces as /ɦɔnt/, because Dutch does not allow phonological words to end in a voiced obstruent. The mismatches in Section 2.2 are examples of such repairs: in the underlying representation of transcription (2), there are two adjacent |s| segments, because |zʏs| *zus* 'sister' and |sofiː| *Sophie* are still divided by a morpheme boundary; at SF, this boundary has been deleted, and therefore a phonotactic restriction that disallows geminates can apply. Similarly, we assume that underlyingly there is still a final |n| in |vɑn| *van* 'of', because this is most likely the form that has been stored in the mental lexicon; however, the labial place of articulation of the initial |p| in |praːχ| *Praag* 'Prague' causes this |n| to surface as a labial /m/ (as in (2)). In Optimality Theory (Prince and Smolensky 1993/2004; hereafter "OT"), the relation between UF and SF is evaluated by so-called faithfulness constraints (McCarthy and Prince 1995), that aim to preserve the information in the UF and prevent any mismatches such as degemination or place assimilation. In the UF–SF mapping, then, this notion of faithfulness is an example of transparency as applied in FDG research by Hengeveld and Leufkens (2018), and any violations of faithfulness induce opacity in the sense it is used in FDG.

However, the term "opacity" is used differently in the phonological literature, where it is a possible property of a transformational rule, or of an interaction of transformational rules. Such rules have the format A → B / C __ D, where A, B, C and D are phonological features, matrices of features, or contexts such as syllable boundaries; the format means that A changes into B if it follows C and precedes D. A and B can also be empty sets, in the case of epenthesis and deletion, respectively; at least one of the contexts C or D needs to be present. A rule, or an interaction of rules, is said to be opaque when one of three conditions is met: (i) there are surface representations in which A occurs in the context C __ D too; (ii) there are surface representations in which the rule or interaction of rules creates B in a different context than C__D; (iii) there are surface representations in which B occurs in the context C__D, but it has not been created by the rule or interaction of rules.

For instance, in Japanese, a phonological process exists that palatalizes consonants before front high vowels; another process deletes high vowels between voiceless obstruents. The UF |sika|, then, is realized as the SF /ɕka/. This interaction is opaque: the palatalization process introduces an alveolopalatal fricative segment, but the trigger for this segment is subsequently deleted by the vowel deletion process. In this example, the palatalization rule needs to apply first; if the high vowel deletion would occur first, we would get an incorrect (i.e. unattested) output. Both the correct and incorrect orders are shown in (3), where the asterisk indicates an unattested form.

(3)

correct derivation:		**incorrect derivation**:	
sika	*underlying form*	sika	*underlying form*
ɕika	*palatalisation*	ska	*high vowel deletion*
ɕka	*high vowel deletion*	ska	*palatalisation (n/a)*
ɕka	*surface form*	*ska	*surface form*

Another example is found in Canadian English. This variety of English has a phonological rule that raises low vowels before voiceless obstruents ("vowel raising"), and another rule that turns coronal stops into flaps if they occur in between vowels ("intervocalic flapping"). The derivations of the words *writing* and *riding* are given in (4). In the underlying representations, the only difference between the forms is the voicing feature value of the final segment of the root; as a consequence of this difference, *writing* undergoes vowel raising while *riding* does not. Subsequently, the voicing difference is obscured by the intervocalic flapping rule. The only difference between the surface forms lies in the quality of the diphthong; the interaction of rules that leads to the SF of *writing* is opaque, because this SF contains a raised diphthong followed by a voiced segment, while raised vowels normally only occur before voiceless segments.

(4)

derivation of *writing*:		**derivation of *riding***:	
ɹaɪt+ɪŋ	*underlying form*	ɹaɪd+ɪŋ	*underlying form*
ɹʌɪtɪŋ	*vowel raising*	ɹaɪdɪŋ	*vowel raising (n/a)*
ɹʌɪɾɪŋ	*intervocalic flapping*	ɹaɪɾɪŋ	*intervocalic flapping*
ɹʌɪɾɪŋ	*surface form*	ɹaɪɾɪŋ	*surface form*

In both these examples, it is impossible to arrive at the correct SF without an intermediate step: the rules need to apply separately and sequentially. Rule-based approaches allow for as many intermediate representations as necessary; each representation is derived from the last through the application of a transformational rule, as in (3) and (4).

However, in OT, the current mainstream framework in generative phonology, an SF is derived from a UF directly, and intermediate representations do not exist, at least in OT as originated by Prince and Smolensky. To account for the data from (3) and (4) with OT, a somewhat hybrid model needs to be assumed that selects an optimal candidate through a ranking of violable constraints (as in OT), but that allows for intermediate representations (as in pre-OT approaches). Such models are usually referred to as Stratal OT (Bermúdez-Otero 1999; Kiparsky 2000), with a "stratum" being an intermediate representation; at every stratum, the OT grammar (i.e. the set of constraints as well as their ranking) is different. OT formalizations of the acquisition of opaque mappings can be found in McCarthy (1999, 2003), Kiparsky (2000), Bermúdez-Otero (2003), Jarosz (2016), Nazarov and Pater (2017), and Prickett (2019). Learners, both simulated and human, prefer transparent mappings over opaque ones (Ettlinger 2008; Kim 2014; Prickett 2019).

It is debated whether opacity is a synchronic, productive process, or instead a historically motivated phenomenon (and, as such, another case of "historical junk" as mentioned earlier), in which case phonological grammars may not need to be able to account for it. The answer probably needs to be established on a case-by-case basis: Donegan and Stampe (1979) give examples of opaque yet productive processes in English, and Al-Mozainy (1981) argues that opacity in Bedouin Arabic is indeed productive, but Kawahara (2017) discusses several kinds of opacity in Japanese for which the evidence for their productivity is mixed. Sanders (2003) divides known cases of opacity into three groups: cases that are synchronically unproductive, cases that are synchronically productive but morphologically conditioned, and cases that can be reanalysed transparently.

In summary, the term "opacity" as used in the phonological literature has a more restricted application than in the FDG literature. In the remainder of this paper, we will restrict the discussion to opacity in the FDG sense of the word.

4 Phonology and phonetics in FDG

As outlined in the previous section, many phonologists assume that at least two levels of representation are needed. In the layout of FDG as presented in Hengeveld and Mackenzie (2008), the model has a single level of phonological representation, which is structured in terms of prosodic constituents and therefore seems to be identical to a traditional surface form. O'Neill (2013) proposed to add an underlying level, and the terminology in Hengeveld and Leufkens (2018: 158) suggests the authors' acknowledgment of this distinction: "purely phonological rules [may] apply that adapt an underlying phoneme to its surface environment". Hengeveld and Mackenzie (this volume) do not explicitly represent the distinction in their model, but still refer to an underlying phonological representation, suggesting that they do in fact recognize the distinction.

As the term "surface form" suggests, many phonologists consider this representation to be the one that is accessible to inspection, i.e. directly measurable: this means either that the SF is translated into a phonetic representation through a universal phonetic encoder that is of no interest to the linguist (the stance taken by Chomsky and Halle 1968, and many other generative phonologists), or that the SF itself contains phonetic detail, specifying a plan of auditory cues and articulatory gestures in addition to phonological prosodic content. Kirchner (1998/2001), for instance, evaluates SF in terms of its articulatory effort. In this view, the SF is a mix of discrete (i.e. phonological) and continuous (i.e. phonetic) properties.

A theoretical framework that teases apart phonological and phonetic representations is Boersma's bidirectional model of phonology and phonetics (cf. Boersma 2011 for an overview). This framework assumes continuous auditory-phonetic and articulatory-phonetic representations in addition to the traditional discrete phonological UF and SF. This separation allows for an explicit formalization of the phonology-phonetics interface. It is also advantageous when explaining a number of phenomena in natural language, such as loanword adaptation (Boersma and Hamann 2009), auditory dispersion (Boersma and Hamann 2008; Seinhorst, Boersma and Hamann 2019) and *h aspiré* in French (Boersma 2007). The BiPhon model is explicitly bidirectional, meaning that it assumes speaker–listeners to use the same knowledge both in perception and in production; FDG focuses on the production direction, but can be used in both directions of processing. The BiPhon and FDG models differ with respect to their stance on parallel/serial processing: in the BiPhon model, "later" representations may influence earlier ones, while FDG assumes strict seriality. (For a more detailed comparison of the two models, see Seinhorst 2014.) We challenge the assumption of seriality in Section 6.

Following the BiPhon's model separation of phonology and phonetics, Seinhorst (2014) proposed to extend FDG with a Phonetic Level. Since the (sub)levels within the Grammatical Component are considered to be discrete representations, this level should be placed in the Output Component, where non-discrete processes take place.

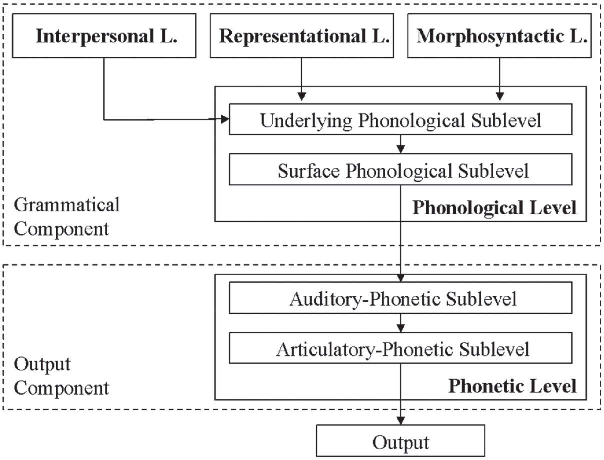

Figure 1: The combined architecture of FDG and Boersma's BiPhon model, proposed by O'Neill (2013), who divided the Phonological Level into two representations, and Seinhorst (2014), who added the Phonetic Level. We further refine this proposal below and in Figure 2.

Figure 1 shows part of the architecture of FDG with the four proposed sublevels, with strictly top-down processing in the production direction (which we will amend in Section 6), and disregarding the Contextual Component for now (but again see Section 6 for a discussion on the role of this component in language change). The Auditory-Phonetic Sublevel is a representation of auditory events, such as formants, frication noise, plosive release bursts, pitches, and so on; the Articulatory-Phonetic Sublevel contains a gestural plan specifying the muscle movements that are needed to produce the utterance.

These phonetic representations are necessary to make FDG compatible with the tenets of functionalist phonology (discussed in Section 2.2). The pressure towards perceptual clarity crucially involves the selection at the Auditory-Phonetic Sublevel of the least ambiguous auditory cues for a phonological input structure, and the pressure towards articulatory ease crucially involves an evaluation of the gestural effort needed to produce the content of the Articulatory-Phonetic Sublevel.

Figure 1 reflects the UF–SF distinction familiar from generative phonology, but as García Velasco (p.c.) pointed out, the architecture of FDG may make it unnecessary to maintain this distinction. In generative phonology, the underlying form is the representation where the phonemic form of an utterance is retrieved from the lexicon once its morphosyntactic shape has been defined. In FDG, by contrast, such phonemic forms may already become available after the Interpersonal Level, if the Representational and Morphosyntactic Levels are skipped (Hengeveld and Mackenzie 2008: 13), or at the Representational Level if the Morphosyntactic Level is skipped (ibid.). In such cases, the output of the formulation levels is sent to the Phonological Encoding process, which we could therefore interpret as the retrieval of morphophonemic forms from the set of lexemes and other primitives, the concatenation of these morphophonemic forms into utterances, and the translation of this concatenation to the prosodic representation that is the Phonological Level. We are agnostic with respect to the status of an underlying phonological sublevel in FDG, and with respect to the question of whether the interface between a lexical representation and a prosodic one is best modelled in FDG as an interaction between separate levels or as part of Phonological Encoding; in the remainder of this paper, we will forgo the assumption of an underlying sublevel, adhering to the architecture proposed by Hengeveld and Mackenzie (2008). The mismatches discussed in this paper do not bear on this question.

We might expect that the effects of the tendency towards articulatory ease are not language-specific, since all humans have (roughly) the same speech apparatus at their disposal, and the notion of effort may therefore be universal. However, languages differ in their use of the available auditory and articulatory space. For instance, a vowel system with three vowel heights will occupy a larger auditory space than a vowel system with only two heights, and the production of the contrasts in this bigger system will require more articulatory effort. At the same time, the advantage of this increased effort is a larger number of contrasts, and hence a lower degree of perceptual confusion. Another example of the language-specificity of the conflict between perceptual clarity and articulatory ease is nasal place assimilation. In many languages, coronal nasals undergo place assimilation to a following consonant: for instance, in example (1), the |n| of |vɑn| 'of' assimilated to /m/ under the influence of the following |p|, deleting the tongue tip gesture that is needed to produce an [n]. In Limburgish, however, coronal nasals remain coronal, maintaining the tip tongue gesture (Boersma 1998: 469). Since the outcome of this conflict is language-specific, we argue that both forces should be reflected somewhere in the Grammatical Component; this will be the topic of Section 6.

5 Interfaces and mismatches between the phonological and phonetic levels of FDG

Within FDG, phonological opacity (as defined in Section 3) would emerge in the Phonological Encoding process, as a result of two consecutive mismatches that should occur in a specific order: in both examples (3) and (4), a substitution is followed by a deletion or by another substitution. However, more interfaces involving phonology and phonetics exist, and therefore many more mismatches occur, giving rise to FDG opacity (i.e. the definition from Section 2.1). This section discusses those interfaces as well as mismatches that may obtain there. We can only discuss a small number of mismatches here, although countless examples are available: an immense degree of variation exists between languages with regard to their phoneme inventories, phonotactic constraints, phonological processes, and phonetic implementation.

As noted in Section 4, information can be taken to the phonology from any higher level (Hengeveld and Mackenzie 2008: 13), but this will usually be the Morphosyntactic Level.

5.1 Mismatches between the morphosyntactic and phonological levels

Hengeveld and Leufkens (2018) provide a number of examples of mismatches between morphosyntactic and phonological units of representation, for instance when morpheme and syllable boundaries do not coincide, or when multiple syntactic elements are expressed within a single morpheme, as in stem or affix alternations: for instance, in Dutch, the verb *lop-en* 'walk-INF' is realized in the past tense as the monomorphemic form |li:p| 'walk.PST.SG'.

Many other opaque phenomena occur at the interface between the Morphosyntactic and Phonological Levels as well. One such phenomenon is phonologically conditioned allomorph selection, such as the choice between *vieux* and *vieil* 'old-M' in French: this choice depends on the following segment, with *vieux* occuring before consonants or *h aspiré*. Other such examples are the choice between the articles *a* and *an* in English, or in some cases, the choice between the articles *el* and *la* in Spanish: even though the Spanish word *agua* 'water' is feminine, it is preceded by the masculine definite article to avoid hiatus. These phenomena all represent one-to-many relations between morphemes and their phonological realisation.

Another example of a mismatch between the Morphosyntactic and Phonological Levels comes from Yawelmani, an almost extinct Yokutsan language. Yawelmani has an underlying contrast between long and short vowels, but long vowels are not allowed to surface before a syllable coda (Kenstowicz and Kisseberth 1979), as can be seen in the examples in (5). In (5a) the verbal root |xil| 'to tangle' has a short vowel, which surfaces as short irrespective of the occurrence of a syllable coda. In (5b), the long vowel in |sa:p| 'to burn' surfaces unaltered in the dubitative form, because the following /p/ belongs to the second syllable. It is, however, shortened in the future passive form, because it precedes a coda consonant.

(5) a. Underlyingly short vowels remain short before a syllable coda:

tangle-DUB	tangle-FUT.PASS	*meaning*
xil+al	xil+nit	*lexical form*
xi.lal	xil.nit	*surface form*

b. Underlyingly long vowels become short before a syllable coda:

burn-DUB	burn-FUT.PASS	*meaning*
sa:p+al	sa:p+nit	*lexical form*
sa:.pal	sap.nit	*surface form*

This last example also shows a different mismatch, already mentioned by Hengeveld and Leufkens (2018): the misalignment of morphophonemic and prosodic units, as a consequence of their fundamentally different natures. The root 'to burn' is |sa:p|, so the underlying form of the dubitative has a morpheme boundary after the |p|; however, the Maximum Onset Principle requires that in the syllabification process, onsets be filled first, so in the surface representation, the /p/ is assigned to the onset of the second syllable rather than to the coda of the first one. Examples of this mismatch abound cross-linguistically; the Maximum Onset Principle is likely rooted in properties of human audition, preferring large sonority contrasts between syllable onsets and syllable nuclei (Delgutte 1982).

Interestingly, many languages display phonological sensitivity to morphological class. In Chuukese, an Austronesian language, nouns need to contain at least two moras, but verbs may surface as monomoraic (Smith 2011), and in Arabic, roots but not affixes may contain pharyngeal segments (McCarthy and Prince 1995). Reduplication, the copying of (part of) the phonological structure of a morphological unit for morphosyntactic purposes, is a cross-linguistically frequent phenomenon as well, often used to pluralize or intensify. Consider, for instance, example (6) from Etsako, a language spoken in Nigeria (Elimelech 1978), in which reduplication of a noun signifies the meaning 'every'.

(6) [ówà] 'house'
[ówǒwà] 'every house'

These phenomena are not examples of mismatches, but they show how morphology and phonology may interact.

5.2 Mismatches between any higher level and the phonological level

Although the Morphosyntactic Level is the most frequent supplier of input to the Phonological Encoding module, it is not the only one. Regardless of the input level, Phonological Encoding translates a phonemic representation into a representation that is structured in terms of intonational and phonological phrases, words, syllables, segments, and possibly moras, if those are needed in the description of the language; this representation obeys language-specific phonotactic restrictions. Hengeveld and Leufkens (2018: 158) already mention a number of phonological processes that may decrease transparency during Phonological Encoding, such as vowel harmony, nasal (place) assimilation, and final devoicing (the process that requires Dutch word-final obstruents to be voiceless, cf. Section 3). Hengeveld (2011) also mentions Dutch degemination (cf. Section 2.2), Spanish diphthongization, and Turkish vowel harmony as opaque features that emerge "when phonological rules apply that adapt an underlying phoneme to its phonological environment", hence on the interface between a phonemic form and the Phonological Level. Leufkens (2011) adds nasalization, segment epenthesis and deletion. Yet other processes that may occur in this interface are, for instance, vowel reduction, tone spreading, tonal sandhi, and so on. These are all examples of phenomena where a mismatch is found between the stored form of the lexeme and the way in which it surfaces in the Phonological Level, violating the notion of faithfulness (cf. Section 3) in order to ensure that the surface representation meets phonotactic requirements.

Example (6) above, repeated here as (7), is not only a case of reduplication but also an example of tone spreading. Etsako has two lexical tones: high (H) and low (L). It also has a phonological rule that deletes a vowel if it is followed by another vowel. In a phrase like |ówà#ówà| 'every house' (example 7a.), for instance, the |à| from the first word needs to be deleted, in order to resolve possible hiatus (i.e. two adjacent vowels in successive syllables). However, the deletion of this vowel leaves its high tone stranded (7b.), since tone is autosegmental, i.e. represented on a tier separately from the segments (Goldsmith 1976; Clements 1976; McCarthy 1981). The stranded tone is then associated with the next syllable. Since this syllable already carries a high tone in addition to the formerly stranded tone, it surfaces with a rising (LH) tone /ówǒwà/ (7c.).

```
(7)  a.  lexical representation          owa   #   owa   (CV tier)
                                         | |       | |
                                         H L       H L   (tonal tier)

     b.  hiatus resolution               ow    #   owa
                                         |         | |
                                         H L       H L

     c.  prosodic representation (PL)    ow    #   owa
                                         |        /| |
                                         H       L H L
```

This process yields a form that obeys the phonotactic restrictions of Etsako, in terms of both segments and tones; note that, although the hiatus resolution induces opacity at the segmental tier because it deletes a vowel, the tone of this vowel was maintained. Thus, the mapping from the lexical representation to the prosodic representation is transparent as far as the tonal tier is concerned.

5.3 The interface between the phonological and phonetic levels

The interface between the Phonological and Phonetic Levels is the interface between the Grammatical Component and the Output Component: this is where the discrete is translated into the continuous, in a process we would like to call

Phonetic Encoding. Following Boersma (2009, 2011), we assume that the prosodic representation (the traditional surface representation, and the Phonological Level in FDG) is translated into an auditory and an articulatory process in parallel, as opposed to the strict top-down processing in Figure 1 above. The Phonetic Encoding process uses all knowledge of the relation between discrete phonological units and their auditory and articulatory correlates, for instance the knowledge that a plosive segment is usually marked by an auditory release, or that a phonologically high vowel has a low first formant, or that an extremely high first formant is articulatory extremely effortful.

Since the phonetic representations are continuous and infinitely variable, no numerical mismatches occur in this interface (remember our definition from Section 2.1): it is hard to imagine what would constitute a match or a mismatch, although speaker–listeners do classify certain tokens as more prototypical instances of a phonological category than others (Johnson, Flemming and Wright 1993; Frieda, Walley, Flege and Sloane 2000).

5.4 The interface between the auditory-phonetic and articulatory-phonetic sublevels

We assume the auditory and articulatory representations to be computed in parallel, which entails that we do not assume the articulatory form to be derived from the auditory form. Nevertheless, speaker–listeners possess sensorimotor knowledge, that is, knowledge of the relation between auditory events and articulatory gestures. An example of sensorimotor knowledge would be that muscles involving the lowering of the jaw need to be active if a vowel with a high first formant is produced. The acquisition of this knowledge already begins when an infant starts babbling, and it may need to be reorganized at any point in the speaker–listener's life, for instance if she has an (innate or acquired) speech impediment that she aims to resolve.

In Figure 1 above, an arrow connects the Articulatory-Phonetic Sublevel and the Output, suggesting the existence of another interface. However, by "Output" we mean the sound waves that eventually impinge on the listener's ear; the transition from articulation to air pressure differences is inherent to the articulation process. It is this Output that feeds into the Contextual Component.

5.5 Mismatches and learnability

One might argue that all mismatches within and between the phonological and phonetic interfaces are equal, but from the point of view of the listener/learner, this may not be true. In the same way that many-to-one relations have different repercussions for learnability than null-to-one-relations like grammatical gender (as argued in Section 2.1), different forms of opacity at the phonological level may have different effects as well. A case in point is that phonological processes can cause neutralization, meaning that they may cause an underlying contrast to be obscured at the surface: in German, final devoicing causes the UF |ʁaːt| *Rat* 'council' and the UF |ʁaːd| *Rad* 'wheel' both to surface as /ʁaːt/. In perception and learning, processing a neutralizing mismatch likely causes more difficulty than a non-neutralizing one. For instance, a Dutch listener, upon perceiving the surface form /ɦɔnt/, does not have to disambiguate between |ɦɔnd| and |ɦɔnt|, as the latter form does not exist. In a sense, neutralization can be regarded as a mismatch on its own, in which one representation at the Phonological Level corresponds to two lexical entries in the Fund, with different semantic and potentially different morphosyntactic properties. While neutralization only occurs in certain prosodic contexts (in this example, the final devoicing process that causes the neutralization between *Rad* and *Rat* only happens at the end of the phonological word), full homonyms would be another example of opaque relationships between items in the lexicon and their semantic and morphosyntactic behaviour, whose phonological representations are already indistinguishable in the Fund.

6 The direction of processing: Bottom-up phonetic influences

It is noteworthy that only three opaque traits are shared by all thirty languages in Hengeveld and Leufkens' (2018) sample: apposition, cross-reference, and phonological alternations. Hengeveld and Leufkens argue that the former two are examples of repetition of information, motivated by a need for expressivity, and that the latter is motivated by a need for articulatory efficiency. We would like to point out that these two motivations are fundamentally very different. The need for expressivity only adds information, while the need for efficiency will only change or delete information, which might hamper communicative success as well as the acquisition process (cf. Section 2). For this reason, the ubiquity of phonological alternations in Hengeveld and Leufkens' sample actually seems much more surprising than the occurrence of the two redundant traits.

Such efficiency, counterproductive as it may be, is motivated by the force of articulatory ease (while taking perceptual clarity into account), so within FDG, it involves the representations within the Phonetic Level. The observation that such efficiency has become obligatory, as part of the phonological grammar, indicates that the Phonetic Level must interface with the Phonological Level somehow, i.e. bottom-up, serving as a bottleneck: auditory and articulatory considerations can force categorical processes to occur, and cause loss of phonological substance, such as syllables, tones, and words. Examples of this bottom-up influence are not only found in reduced speech, but also in the stepwise process of grammaticalization. Grammaticalizing constructions usually undergo a rise in frequency, which makes them more predictable: the speaker will therefore be more likely to expend less articulatory effort to produce the construction, making it more prone to phonological reduction and erosion (Bybee, Perkins, and Pagliuca 1993; Bybee 2003; Hopper and Traugott 2003; within FDG: Keizer 2007; Olbertz 2007; Grández Ávila 2010). Probably the most well-studied instance of grammaticalization is the English periphrastic future marker *to be going to*: *I'm going to* is commonly reduced to *I'm gonna* or even *imma* (a.o. Givón 1979; Bybee, Perkins, and Pagliuca 1994; Hopper and Traugott 2003). In this last form, only two syllables remain. In a survey of the use of *don't* in spoken American English, Scheibman (2000) found that *don't* was reduced most, to [ə], after the pronoun *I* and before high-frequency verbs such as *know*. Not only does this reduction display significant segmental deletion, it has also lost its prosodic independence: it is no longer a phonological word, but a clitic.

Evidence for the language-specificity of this bottom-up influence is provided, for instance, by Wanrooij and Raijmakers (2020, under review), who show that languages have their own reduction rules: German and Dutch infinitive verbs both end in |ən|, but under reduction, German infinitives undergo schwa deletion and consecutive assimilation of consonants (e.g. *haben* |haːbən| 'to have' is reduced from haː.bən to haːb.n > haːb.m > haːm.m > haːm (Kohler 1996)) whereas in Standard Dutch *hebben* the word-final /n/ is deleted. We interpret this observation as evidence that the influence of phonetics on phonology must have a place in the Grammatical Component: the phonetic process of reduction is grammaticalized into language-specific, obligatory phonological rules. Importantly, the Dutch and German rules differ for good reason. If the word-final /n/ were dropped from the German infinitive |haːbən| (the "Dutch" strategy), we would get the form /haːbə/; applying the "German" strategy to the Dutch infinitive |hɛbən| would yield /hɛm/. Both of these words already exist in the language: in German, /haːbə/ is the form of the first person singular of the same verb, and in Dutch, /hɛm/ is the personal pronoun "him". Both languages, then, show phonetically-based reduction, but they do so in a way that minimizes confusion with other lexical items. Hengeveld and Mackenzie (this volume) argue that phonetic processes such as reduction, assimilation and degemination occur in the Articulator, and that the reduced form may eventually become available as a lexical entry through a feedback loop involving the Contextual Component. While we agree that the Contextual Component is indeed relevant in this diachronic entrenchment, the Articulator is not part of the

Grammatical Component, and therefore Hengeveld and Mackenzie's solution fails to account for the German and Dutch data by denying that reduction proceeds in a language-specific way which takes lexical considerations into account.[3]

We model the effect of auditory and articulatory considerations on a prosodic representation as a bottom-up influence of the Phonetic Encoding process on the Phonological Encoding process, which we believe best explains that a need for articulatory ease, as far as it is allowed by lexical and perceptual factors, may result in a phonologically reduced representation. Figure 2 shows a modified version of Figure 1, now including the bottom-up interface by an arrow from the Phonetic Encoding process to the Phonological Encoding process. The grey double-sided arrow between the Auditory- and Articulatory-Phonetic Sublevels indicates that the latter is not derived unidirectionally from the former, but that speaker–listeners do possess knowledge of the relation between these two forms.

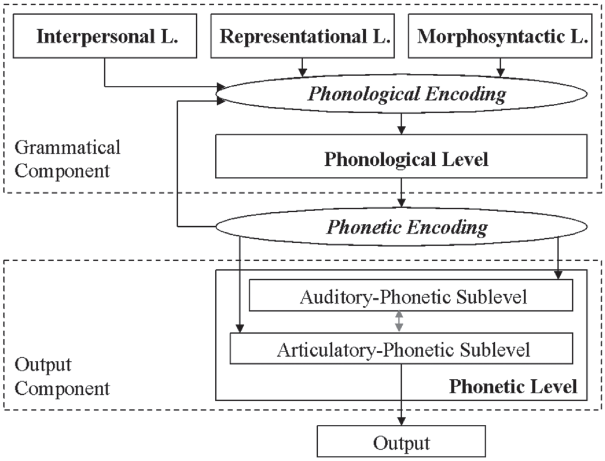

Figure 2: The coupling of FDG and the BiPhon model from Figure 1, as proposed by O'Neill (2013) and Seinhorst (2014), now including a bottom-up feedback loop from the phonetics to the phonology.

Although we place the Phonetic Encoding process outside the Grammatical Component because it involves non-discrete representations, we would like to stress once more that this process is language-specific, since it is fed by the Phonological Level, and since the weighing of lexical, perceptual and artic-

3 As we mentioned in Sections 2.2 and 4, assimilation processes are language-specific too, and so is degemination. We therefore argue that these processes do not belong in the Articulator, but rather in Phonological Encoding.

ulatory factors is language-specific. We believe that the introduction of this bottom-up bottleneck allows FDG to increase its explanatory adequacy, and strengthen its functional nature. We would also like to note that, even though the phonological processes and examples that are described in this paper are taken from spoken languages, the basic notions extend to sign language as well, which is equally subject to the forces of perceptual clarity and articulatory ease (a.o. Crasborn 2001; for grammaticalization in sign languages, cf. Pfau and Steinbach 2011).

7 Conclusion

In this paper we have discussed extensions to the FDG model's architecture that were proposed by O'Neill (2013) and Seinhorst (2014), as well as some mismatches that may occur between the phonological and phonetic levels of FDG. These processes seem to be motivated by perceptual and articulatory considerations, pointing to a bottom-up influence from the phonetics on the phonology. We have assumed that mismatches can only occur between discrete units of representation, i.e. that mismatches may only occur with and within the Phonological Level; since numerous phonological processes can cause opacity, it will be interesting to compare the types of mismatches at the Phonological Level with those found at the other levels of representation. Finally, we have argued that the tenets of functional phonology, which state that language users aim to speak as efficiently as they can while still being understood correctly, can be incorporated into FDG through the introduction of a bottom-up influence from the phonetics to the phonology, increasing the model's explanatory adequacy.

References

Al-Mozainy, Hamza. 1981. *Vowel alternations in a Bedouin Hijazi Arabic Dialect: abstractness and stress*. Austin, TX: University of Texas at Austin dissertation.

Audring, Jenny. 2014. Gender as a complex feature. *Language Sciences* 43. 5–17.

Barbiers, Sjef. 2008. Microvariation in syntactic doubling – an introduction. In Sjef Barbiers, Olaf Koeneman, Marika Lekakou & Margreet van der Ham (eds.), *Microvariation in syntactic doubling*, 1–34. Leiden: Brill.

Bermúdez-Otero, Ricardo. 1999. *Constraint interaction in language change: quantity in English and Germanic*. Manchester: University of Manchester dissertation.

Bermúdez-Otero, Ricardo. 2003. The acquisition of phonological opacity. In Jennifer Spenader, Anders Eriksson & Östen Dahl (eds.), *Variation within Optimality Theory: Proceedings of the Stockholm Workshop on Variation within Optimality Theory*, 25–36. Stockholm: Department of Linguistics, Stockholm University.
Blom, Elma, Daniela Polišenská & Fred Weerman. 2008. Articles, adjectives and age of onset: the acquisition of Dutch grammatical gender. *Second Language Research* 24. 297–331.
Boersma, Paul. 1998. *Functional Phonology: formalizing the interactions between articulatory and perceptual drives*. Amsterdam: University of Amsterdam dissertation.
Boersma, Paul. 2007. Some listener-oriented accounts of *h-aspiré* in French. *Lingua* 117. 1989–2054.
Boersma, Paul. 2009. Cue constraints and their interactions in phonological perception and production. In Paul Boersma & Silke Hamann (eds.), *Phonology in perception*, 55–110. Berlin & New York: Mouton de Gruyter.
Boersma, Paul. 2011. A programme for bidirectional phonology and phonetics and their acquisition and evolution. In Anton Benz & Jason Mattausch (eds.), *Bidirectional Optimality Theory*, 33–72. Amsterdam: John Benjamins.
Boersma, Paul & Silke Hamann. 2008. The evolution of auditory dispersion in bidirectional constraint grammars. *Phonology* 25. 217–270.
Boersma, Paul & Silke Hamann. 2009. Loanword adaptation as first-language phonological perception. In Andrea Calabrese & W. Leo Wetzels (eds.), *Loanword phonology*, 11–58. Amsterdam: John Benjamins.
Bybee, Joan. 2003. Mechanisms of change in grammaticization: the role of frequency. In Richard Janda & Brian Joseph (eds.), *Handbook of historical linguistics*, 602–623. Oxford: Blackwell.
Bybee, Joan, Revere Perkins & William Pagliuca. 1994. *The evolution of grammar: tense, aspect and modality in the languages of the world*. Chicago: Chicago University Press.
Chomsky, Noam & Morris Halle. 1968. *The sound pattern of English*. New York: Harper and Row.
Clements, G. Nick. 1976. *Vowel harmony in nonlinear generative phonology: an autosegmental model*. Indiana University Linguistics Club.
Coles-White, D'Jaris. 2004. Negative concord in child African American English: implications for Specific Language Impairment. *Journal of Speech, Language, and Hearing Research*, *47*. 212–222.
Contreras-García, Lucia. 2013. *Grammar in 3D: On linguistics theory design*. Amsterdam: University of Amsterdam dissertation.
Crasborn, Onno. 2001. *Phonetic implementation of phonological categories in Sign Language of the Netherlands*. Leiden: Leiden University dissertation.
Ćurčić, Maja. 2018. *Explaining differences in adult second language learning: the role of language input characteristics and learners' cognitive aptitudes*. Amsterdam: University of Amsterdam dissertation.
Dahl, Östen. 2004. *The growth and maintenance of linguistic complexity*. Amsterdam & Philadelphia: John Benjamins Publishing Company.
De Houwer, Annick & Steven Gillis. 1998. Dutch child language: an overview. In Steven Gillis & Annick De Houwer (eds.), *The acquisition of Dutch*, 1–100. Amsterdam: John Benjamins.
Delgutte, Bertrand. 1982. Some correlates of phonetic distinctions at the level of the auditory nerve. In Rolf Carlson & Björn Granström (eds.), *The representation of speech in the peripheral auditory system*, 131–150. Amsterdam: Elsevier.

Donegan, Patricia & David Stampe. 1979. The study of Natural Phonology. In Daniel Dinnsen (ed.), *Current approaches to phonological theory*, 126–173. Bloomington, IN: Indiana University Press.

Elimelech, Baruch. 1978. *A tonal grammar of Etsako*. Berkeley: University of California Press.

Ernestus, Mirjam. 2000. *Voice assimilation and segment reduction in casual Dutch: a corpus-based study of the phonology–phonetics interface*. Amsterdam: Free University dissertation.

Ettlinger, Marc. 2008. *Input-driven opacity*. Berkeley, CA: University of California at Berkeley dissertation.

Frieda, Elaina, Amanda Walley, James Flege & Michael Sloane. 2000. Adults' perception and production of the English vowel /i/. *Journal of Speech, Language, and Hearing Research* 43. 129–143.

Givón, Talmy. 1979. *On understanding grammar*. New York, NY: Academic Press.

Goldsmith, John. 1976. *Autosegmental phonology*. Cambridge, MA: Massachusetts Institute of Technology dissertation.

Grández Ávila, Magaly. 2010. *A functional approach to the subjectification of facultative meaning: the case of* capaz *in American Spanish*. Amsterdam: University of Amsterdam MA thesis.

Hendriks, Petra & Helen de Hoop. 2001. Optimality-Theoretic semantics. *Linguistics and Philosophy* 24. 1–32.

Hengeveld, Kees (ed.). 2011. Transparency in Functional Discourse Grammar. *Linguistics in Amsterdam* 4.

Hengeveld, Kees & Sterre Leufkens. 2018. Transparent and non-transparent languages. *Folia Linguistica* 52. 139–175.

Hengeveld, Kees & J. Lachlan Mackenzie. 2008. *Functional Discourse Grammar: a typologically-based theory of language structure*. Oxford: Oxford University Press.

Hengeveld, Kees & J. Lachlan Mackenzie. This volume. *Interfaces, mismatches, and the architecture of Functional Discourse Grammar*.

Hopper, Paul & Elizabeth Traugott. 2003. *Grammaticalization*. Cambridge: Cambridge University Press.

Hudson Kam, Sharon & Elissa Newport. 2005. Regularizing unpredictable variation: the roles of adult and child learners in language formation and change. *Language Learning and Development* 1. 151–195.

Jarosz, Gaja. 2016. Learning opaque and transparent interactions in Harmonic Serialism. In Gunnar Ólafur Hansson, Ashley Farris-Trimble, Kevin McMullin & Douglas Pulleyblank (eds.), *Proceedings of the 2015 Annual Meeting on Phonology*.

Johnson, Keith. 2004. Massive reduction in conversational American English. In Kiyoko Yoneyama & Kikuo Maekawa (eds.), *Spontaneous speech: data and analysis. Proceedings of the 1st Session of the 10th International Symposium*, 29–54. Tokyo: The National International Institute for Japanese Language.

Johnson, Keith, Edward Flemming & Richard Wright. 1993. The hyperspace effect: phonetic targets are hyperarticulated. *Language* 69. 505–528.

Kahn, Daniel. 1976. *Syllable-based generalizations in English phonology*. Cambridge, MA: Massachusetts Institute of Technology dissertation.

Kawahara, Shigeto. 2017. *Phonological opacity in Japanese*. Unpublished manuscript, Keio University, Tokyo, Japan.

Keizer, Evelien. 2007. The grammatical-lexical distinction in Functional Discourse Grammar. *Alfa – Revista de Lingüística* 51 (special issue: *Advances in Functional Discourse Grammar*). 35–56.

Kemps, Rachel, Mirjam Ernestus, Robert Schreuder & R. Harald Baayen. 2004. Processing reduced word forms: the suffix restoration effect. *Brain and Language* 19. 117–127.

Kenstowicz, Michael & Charles Kisseberth. 1979. *Generative phonology: description and theory*. New York: Academic Press.

Kim, Yun Jung. 2014. Do learners prefer transparent rule ordering? An artificial language learning study. In Andrea Beltrama, Tasos Chatzikonstantinou, Jackson L. Lee, Mike Pham & Diane Rak (eds.), *Papers from the 48th Meeting of the Chicago Linguistic Society*, 375–386. Chicago, IL: Chicago Linguistic Society.

Kiparsky, Paul. 1973. Abstractness, opacity and global rules. In Osamu Fujimura, (ed.), *The dimensions of linguistic theory*, 57–86. Tokyo: Tokyo Institute for Advanced Studies of Language.

Kiparsky, Paul. 2000. Opacity and cyclicity. *The Linguistic Review* 17. 351–367.

Kirchner, Robert. 1998/2001. *An effort-based approach to consonant lenition*. Los Angeles, CA: University of California at Los Angeles dissertation. Rutgers Optimality Archive 276. Published by Routledge (London and New York).

Kohler, Klaus. 1998. The disappearance of words in connected speech. *ZAS Working Papers in Linguistics* 11. 21–34.

Koopmans-van Beinum, Florien. 1980. *Vowel contrast reduction: an acoustic and perceptual study of Dutch vowels in various speech conditions*. Amsterdam: University of Amsterdam dissertation.

Kusters, Wouter. 2003. *Linguistic complexity: The influence of social change on verbal inflection*. Leiden: Leiden University dissertation.

Lass, Roger. 1997. *Historical linguistics and language change*. Cambridge: Cambridge University Press.

Lieberman, Erez, Jean-Baptiste Michel, Joe Jackson, Tina Tang & Martin Nowak. 2007. Quantifying the evolutionary dynamics of language. *Nature* 449. 713–716.

Leufkens, Sterre. 2015. *Transparency in language*. Amsterdam: University of Amsterdam dissertation.

Leufkens, Sterre. To appear. A functionalist typology of redundancy. To appear in *Revista da ABRALIN: Special issue on Functional description of languages*.

Liljencrants, Johan & Björn Lindblom. 1972. Numerical simulation of vowel quality systems: the role of perceptual contrast. *Linguistics* 48. 839–862.

Lupyan, Gary & Rick Dale. 2010. Language structure is partly determined by social structure. *PloS one* 5(1). e8559.

Martinet, André. 1960. *Éléments de linguistique générale*. Paris: Armand-Collin.

McCarthy, John. 1981. A prosodic theory of non-concatenative morphology. *Linguistic Inquiry* 12. 373–418.

McCarthy, John. 1999. Sympathy and phonological opacity. *Phonology* 16. 331–399.

McCarthy, John. 2003. Sympathy, cumulativity, and the Duke-of-York gambit. In Caroline Féry & Ruben van de Vijver (eds.), *The syllable in Optimality Theory*, 23–76. Cambridge: Cambridge University Press.

McCarthy, John & Alan Prince. 1995. Faithfulness and reduplicative identity. In Jill Beckman, Laura Walsh Dickey & Suzanne Urbanczyk (eds.), *Papers in Optimality Theory*, 249–384.

University of Massachusetts Occasional Papers 18. Amherst, MA: Graduate Linguistic Student Association.

Nazarov, Aleksei & Joe Pater. 2017. Learning opacity in Stratal Maximum Entropy Grammar. *Phonology* 34. 299–324.

Nichols, Johanna. 2009. Linguistic complexity: a comprehensive definition and survey. In Geoffrey Sampson, David Gil & Peter Trudgill (eds.), *Language complexity as an evolving variable*, 110–125. Oxford: Oxford University Press.

Olbertz, Hella. 2007. 'Dizque' in Mexican Spanish: the subjectification of reportative meaning. *Rivista di Linguistica* 19 (1). 151–172.

O'Neill, Gareth. 2013. *Initial consonant mutation in Irish Gaelic: a Functional Discourse Grammar analysis*. Amsterdam: University of Amsterdam MA thesis.

Padgett, Jaye. 2003. Contrast and post-velar fronting in Russian. *Natural Language and Linguistic Theory* 21. 39–87.

Passy, Paul. 1890. *Étude sur les changements phonétiques et leur caractères généraux*. Paris: Firmin-Didot.

Petré, Peter. 2019. *The interaction between extravagance and competition in syntactic change*. Paper presented at the conference of the Societas Linguistica Europaea 52, Leipzig University, August 21, 2019.

Pfau, Roland & Markus Steinbach. 2011. Grammaticalization in sign languages. In Heiko Narrog & Bernd Heine (eds.), *The Oxford handbook of grammaticalization*, 683–695. Oxford: Oxford University Press.

Prickett, Brandon. 2019. Learning biases in opaque interactions. *Phonology* 36. 627–653.

Prince, Alan & Paul Smolensky. 1993/2004. *Optimality Theory: constraint interaction in Generative Grammar*. RuCCS Technical Report 2, Rutgers University, Piscateway, New Jersey: Rutgers University Center for Cognitive Science. Revised version published 2004 by Blackwell.

Sanders, Nathan. 2003. *Opacity and sound change in the Polish lexicon*. Santa Cruz, CA: University of California at Santa Cruz dissertation.

Scheibman, Joanne. 2000. *I dunno*: a usage-based account of the phonological reduction of *don't* in American English conversation. *Journal of Pragmatics* 32. 105–124.

Seinhorst, Klaas. 2014. Phonetics in Functional Discourse Grammar. *Web Papers in FDG* 87.

Seinhorst, Klaas. 2017. *Feature economy vs. logical complexity in phonological pattern learning*. *Language Sciences* 60. 69–79.

Seinhorst, Klaas, Paul Boersma & Silke Hamann. 2019. Iterated distributional and lexicon-driven learning in a symmetric neural network explains the emergence of features and dispersion. *Proceedings of the 19th International Congress of Phonetic Sciences*.

Selkirk, Elizabeth. 1981. *English compounding and the theory of word-structure*. In Michael Moortgat, Harry van der Hulst & Teun Hoekstra (eds.), *The scope of lexical rules*. Dordrecht: Foris.

Smith, Jennifer. 2011. Category-specific effects. In Marc van Oostendorp, Colin Ewen, Beth Hume & Keren Rice (eds.), *Companion to Phonology*, 2439–2463. Malden, MA: Wiley-Blackwell.

Smith, Kenny & Elizabeth Wonnacott. 2010. Eliminating unpredictable variation through iterated learning. *Cognition* 116. 444–449.

ten Bosch, Louis. 1991. *On the structure of vowel systems: aspects of an extended vowel model using effort and contrast*. Amsterdam: University of Amsterdam dissertation.

Trudgill, Peter. 2009. Sociolinguistic typology and complexification. In Geoffrey Sampson, David Gil & Peter Trudgill (eds.), *Language complexity as an evolving variable,* 98–109. Oxford: Oxford University Press.
Trudgill, Peter. 2011. *Sociolinguistic typology: social determinants of linguistic complexity*. Oxford: Oxford University Press.
Velde, Freek van der. 2014. Degeneracy: The maintenance of constructional networks. In Ronny Boogaart, Timothy Colleman & Gijsbert Rutten (eds.), *Extending the scope of construction grammar*, 141–179. Berlin: Mouton de Gruyter.
Velde, Marlies van der. 2004. L'acquisition des déterminants en L1: une étude comparative entre le français et le néerlandais. *Acquisition et Interaction en Langue Etrangère* 21. 9–46.
Wanrooij, Karin & Maartje E.J. Raijmakers. 2020. Evidence for immature perception in adolescents: adults process reduced speech better and faster than 16-year olds. *Language Acquisition* 27. 434–459.
Wanrooij, Karin & Maartje Raijmakers. Under review. "*Hama*"? Reduced pronunciations in non-native natural speech obstruct high-school students' comprehension severely at lower processing levels.
White, Lydia, Elena Valenzuela, Martyna Kozlowska-Macgregor, Ingrid Leung & Hela Ben Ayed. 2001. The status of abstract features in interlanguage: gender and number in L2 Spanish. In Anna Do, Laura Domínguez & Aimee Johansen (eds.), *Proceedings of the 25th annual Boston University Conference on Language Development*, 792–802.