Phonological Opacity as Differential Classification of Sound Events
Aleksei Nazarov

1. Introduction

In this paper, I will introduce a theory of phonological representation named "differential classification" and show that this theory allows the representation of phonological opacity in standard Optimality Theory (Prince & Smolensky 1993/2004, McCarthy & Prince 1994) without any extra assumptions. It has been noted by many that phonological opacity, while being an empirically real phenomenon (see especially McCarthy 2007b), is mostly impossible to account for in standard Optimality Theory when traditional representations are assumed. This has triggered the formulation of various approaches that extend standard Optimality Theory (OT) to opaque patterns (such as Lexical Phonology – Kiparsky 2000, Bermúdez-Otero 1999, or Optimality Theory with Candidate Chains – McCarthy 2007b). However, I argue that assuming differential classification as the representational framework instead of standard representations makes it unnecessary to add such extensions onto standard OT. Furthermore, I will argue that differential classification is to some degree implied in some previous work in phonology and speech perception/production. Finally, I will show that the differential classification theory of opacity makes empirical predictions distinct from those of other theories of opacity.

In the rest of this section, I will introduce the theory of differential classification (1.1), propose an implementation of this theory in OT (1.2), and show how this implementation gives us the tools to account for opaque as well as transparent generalizations (1.3).

1.1 Differential versus uniform classification

Differential classification is a view of phonological representation which maintains that the grammar represents sound events at multiple taxonomic levels, and that representations of the same phonetic token at different taxonomical levels are fundamentally independent from one another. For instance, a segment and the features that belong to it are separate and independent pieces of information encoded in the grammar. This can be contrasted with the canonical mainstream view of phonological representation (for instance, as proposed in Chomsky & Halle 1968), in which the only level of classifying linguistic sounds that is relevant to the phonological grammar is the phonological feature.

I will explain below what is meant by different taxonomical levels in general, but for the purpose of this paper, I will assume only two taxonomically distinct levels which seem reasonably uncontroversial: the segment as a holistic entity (for instance, [a] versus [t] versus [q]), and the phonological feature as standing for classes of segments (for instance, [+vocalic] as standing for all vowel segments). I will not consider the purely acoustic/articulatory role of phonological features: instead, I will concentrate only on the classificatory function of features. Thus, features will be taken to stand for a set of segments, but not necessarily for some phonetic parameter or phonetic property.

Plugging this assumption of segments and features as the two relevant taxonomical levels into the definition of differential classification given above, differential classification means that the grammar represents phonetic tokens both in terms of holistic segments and in terms of phonological features, and these two ways of representing a phonetic event are fundamentally independent of one another. This can be contrasted with the standard view of representation, in which the concepts “segment” and “feature” do co-exist, but in a different way: the identity of segments is fully made up of the features associated with them, so that segments are not independent from the features associated with them. Rather, features are mere attributes to segments, while segments are mere anchors for the timing of features.

Fundamental independence of featural and segmental representational units means that the co-occurrence of particular features and particular segments is fundamentally free: any feature value could theoretically occur in combination with any segment value, just as any two segments are theoretically allowed to co-occur next to each other. This can be clearly contrasted with the mainstream view of phonological representation, according to which a particular segment value must co-occur with the feature values that define it: [e] must be assigned [+ATR], [-high], [-round], and may never be assigned [-ATR], [+high] or [+round].

Because differential classification maintains that there is no inherent prohibition against classification of the same sound event as a set of contradictory labels, the connection between sound events and certain classifiers, as well as the connection between different classifiers, needs to be represented in the grammar of individual languages. As will be seen in section 1.2 below, I propose to represent such connections by means of Optimality Theoretic (and thus, violable) constraints which penalize unwanted combinations of various classifiers. For instance, there may be a constraint “/e/ → [+ATR]”, which penalizes any event simultaneously classified as /e/ and as [-ATR].

The details of the operation of such constraints will be explained later, but the fact that the grammar must regulate the coherence of various phonological categories is essential for differential classification to be a useful theory of representations – because if features and segments had no preferred association patterns between them, phonological grammar would have no effect on pronunciation.

By contrast, the standard view of phonological representations – which could be termed *uniform classification* – holds that all phonological representation can be reduced to bundles of features: segmental categories, for instance, can be derived from the set of features associated with a timing unit, and have no independent status. Thus, there is only one taxonomic level.

This means that, for a certain (phonetic) pronunciation, an Optimality Theoretical grammar with standard representations needs to consider fewer candidates than when differential classification is assumed. For instance, the last consonant in a phonetic form [bad] has the following possible representations in a standard uniform classification framework with optional underspecification of coronal consonants:

(1)
a. (full specification)

```
 ___________[coronal]_

 _____[+voice]____


   [b    a    d]
```

b. (place underspecification)

____

_______[+voice]____

[b a d]

However, in a differential classification framework, many more potential representations must be considered which would be prohibited *a priori* in a uniform classification framework, some of which are shown below:

(2)

a. (“wrong” voice feature)

___________[coronal]_

_____[-voice]____

[b a d]

b. (“wrong” place feature)

___________[dorsal]_

_____[+voice]____

[b a d]

c. (non-specification of voice)

___________[coronal]_

____

[b a d]

d. (no specification at all)

____

____

[b a d]

Returning to the general concept of taxonomic levels, I assume that one unit is at a higher taxonomic level than another unit when the former unit describes more potential phonetic tokens than the latter unit. Thus, if we look at concepts familiar from structuralism, a phoneme unit would be on a higher taxonomic level than any of its allophones, and an archiphoneme would be on a higher level than the phonemes it stands for. Similarly, an individual segment category like [p] would be at a higher taxonomic level than a syllable-size chunk like [pa], because [p] would describe any utterance that includes the segment [p], regardless of whether [p] precedes [a] in the same syllable, while the syllable chunk [pa] would describe only those utterances in which [p] indeed does precede [a] in the same syllable.

Thus, differential classification would claim in general that all phonologically significant units are independent of one another, even when they classify the same event. If there is evidence that allophones and phonemes are to be represented as separate units in grammar, then differential classification would maintain that these units are independent and may potentially mismatch, even when they describe the same pronunciation. The same is the case for segments and syllable chunks, or any other pair of types of units that differ in their taxonomic status. As said before, this paper only assumes segments and features as phonologically relevant units, but this assumption has been made rather arbitrarily: there is no strong evidence that segments and features are the only taxonomic levels of phonological representation. Rather, it is an empirical question at which taxonomic levels there are phonologically relevant units: a question to be resolved through theoretical and experimental work. It may even well be that the taxonomic levels required for the mental representation of a language may be specific to that language, making the question of discovering a taxonomy of phonological and phonetic units a language-specific task. However, all of these questions will not be addressed in the confines of this paper.

One final remark is in order: since differential classification presupposes a taxonomic model of representation, and phonetically detailed representations describe a narrow rather than broad range of potential phonetic tokens, one can say that phonetics is at the very bottom of the taxonomy:

Phonological features

Segments

Phonetically detailed descriptions

Phonetic token 1   *Phonetic token 2*   *Phonetic token 3*

Thus, the lowest taxonomic level considered in an analysis will always be closest to phonetics. Since the current model only includes features and segments, it will be the segments that determine pronunciation, not the features. This is because features have a purely classificatory role in this model: they do not stand for phonetic dimensions, but rather for groups of segments.

Summarizing, differential classification is a theory of representation which hypothesizes that different taxonomic levels of phonological representation (segments and features) are independent, and may mismatch to the extent permitted by the grammar, which deviates from the standard assumption that the identity of a segment is determined by its features. The notion of “mismatch to the extent permitted by the grammar" will be explored in the following section, which discusses an implementation of differential classification in Optimality Theory.

1.2 Differential classification in Optimality Theory

In the previous section, it was explained what kind of representations are entailed by the theory of differential classification: representations with independent functions for units at different levels of taxonomy. In this section, I will lay out an implementation of this idea in Optimality Theory. This implementation will lead into an explanation of how phonological opacity can be handled by such a model, which will be shown in the following section.

In the current implementation, I will use two taxonomic levels of phonological classification which are familiar from mainstream work in phonology: segments and phonological features. Segments are on a lower taxonomic level, as they represent holistic categories with a reasonably well-defined (albeit context-sensitive) translation into phonetic values. Features, however, are on a higher taxonomic level, as they group segments into classes which can be employed in generalizations.

As has already been touched upon in the previous section, the sense in which phonological features are used in this implementation is only one of two senses in which features are used in mainstream phonology. The notion "phonological feature" in mainstream phonological theory also has a sense of denoting some articulatory (or acoustic) parameter. For instance, the consonants which are [+sonorant] are defined by Chomsky & Halle (1968) as being produced "with a vocal tract cavity configuration in which spontaneous voicing is possible".

This phonetic or substantive notion of phonological feature will not be used in the current implementation. Rather, I will use "feature" as a synonym for "label which classifies segments". This is the same approach to phonological features which has been taken by Mielke (2004): features are seen as labels which are induced by language learners to represent classes of segments functioning together in phonological regularities.

Of course, segments and features also coexist in mainstream phonology. However, mainstream does not assign a special status to the concept "segment": segments are merely root nodes or timing slots which are associated with a certain set of features. The current implementation, however, does assign a special status to segments: segments are labels assigned to certain abstracted fragments of the speech stream, while features are labels assigned to certain segments. At the same time, features are not seen as mere "attributes" without an independent status, but are seen as having the same formal status as segments: labels which stand for a certain class of phonetic events. This independent status for segments and features can be formalized in a special kind of autosegmental representation.

I assume that representations are autosegmental in the sense that different aspects of the description of a phonetic event are on separate tiers. For instance, segmental values such as [a] or [s] inhabit a separate segmental tier. Each individual feature will also inhabit its own individual tier. As in standard autosegmental theory, association lines between labels on different tiers indicate that these labels describe the same phonetic (sub)event.

The only substantial difference between standard autosegmental representations and what I am proposing here is the fact that features are associated with root nodes in the standard theory, while the current model has a separate tier for segment values, to which feature values are attached. The difference between the standard model and the current model is that a segmental value and the features attached to it may "mismatch" (in the sense that the features attached to a segment may not be the canonical features for that segment):

(3)
a. "matching" feature-segment combination

------[+voice]-----

---------[b]----------

b. "mismatching" feature-segment combination

------[-voice]-----

---------[b]----------

In the case of a "mismatching" combination of features and segments, the pronunciation is determined by the segmental value, not by the feature value(s) attached to it. This is because I assume that features have no phonetic "substance", but merely label groups of segments. At the same time, segments are assumed to be closer connected to phonetic patterns.

Of course, if features do not determine phonetics, and segmental values are allowed to be combined with any set of features, then features have no function whatsoever. In order for features to have a function, there must be some way of relating features to segments, especially since features have no inherent phonetic value in the current system. I propose that the connection between features and segments is regulated by means of a set of violable constraints, ranked in a language-specific way. Thus, the assignment of features to segments is language-specific, and potentially non-uniform.

If the assignment of feature F to a certain segment S is demanded by a violable constraint, then the grammar can sometimes (in some special condition) pick a different feature, G, to be assigned with segment S – provided that there is a higher-ranked constraint which prefers this. This non-uniform assignment of features to segments will be crucial for accounting for phonological opacity, as will be seen later.

Such non-uniform behavior is only possible when the range of representations to be considered as candidates by the grammar also include non-canonical feature-segment combinations (for instance [d] and [-voice], or [e] and [-ATR]). The easiest way to accomplish this is to allow any feature-segment combination to be present in representations, so that the GEN component of Optimality Theory may suggest any non-canonical feature-segment combination in a candidate. In this way, constraint ranking has maximal influence on feature assignment to segments.

Constraints which regulate the assignment of features to segments are of a different kind than regular

Markedness or Faithfulness constraints in Optimality Theory. Formally speaking, these constraints are Markedness constraints, since they do not refer to the discrepancy between input and output forms, but only evaluate outputs. However, the purpose of such constraints is not so much to rule out certain sounds or sound combinations as to classify segments into featural categories. Perhaps such feature-to-segment constraints could be called “paradigmatic” Markedness constraints, in the sense that they do not restrict configurations of sounds, but restrict the range of functions that each sound can have with respect to other sounds.

I assume that the shape of feature-segment association constraints is that of a logical implication: if there is a segmental value [X], then it should be associated with the featural value [F] – or *vice versa* (if feature [F], then segment [X]). One example of such a constraint could be one which desires for the vowel segment [a] to have the feature [-ATR], which could be formulated in the following way:

[a] → [-ATR]: Assign one violation mark for every instance of [a] which is not associated with the feature value [-ATR].

This constraint is satisfied by the configuration in (5a) below, but violated by the configurations in (5b) (in which [a] either has no value for [±ATR] or is associated with the value [+ATR]):

(4)

a. satisfies [e] → [+ATR]	**b.** violates [e] → [+ATR]
--------------[-ATR]-- m i n a	------------- ------------[+ATR]-- m i n a m i n a

Because constraints like [a] → [-ATR] are violable, candidates which do not obey this constraint (such as the ones in (...b)) may win if some constraint which prefers them is ranked above [a] → [-ATR]. For instance, if a (hypothetical) constraint against [-ATR] features in word-final syllables outranked [a] → [-ATR], then /mina/ would be able to surface without initial [+ATR], but still with the segment [a]. This situation is shown in the tableau below.

(5)

/m i n a/	*Final[-ATR]	[e] → [+ATR]
------------[+ATR]--- m i n a		*
-------------[+ATR]-- m i n a	*!	

Thus, feature-to-segment constraints (or “paradigmatic” Markedness constraints) serve to establish a preferential assignment of a certain segment to a certain feature category, but this preferential assignment may be overridden when higher-ranked constraints prefer something else.

Apart from the type of Markedness constraints which regulate classification of segments into featural categories as discussed above, the current model also makes use of regular Markedness constraints and Faithfulness constraints. However, the kind of representations that I propose also implies slight differences in how these other, more familiar types of constraints function. These deviations are discussed below.

The fact that segmental labels and featural labels are formally independent means that some constraints can target only segments, other constraints can target only features, and yet others can target both segments and features. Thus, constraints can target different levels of abstractness. For "regular" Markedness constraints, this means that it is more difficult to determine whether two constraints conflict with one another. When two Markedness constraints refer to distinct levels of abstractness (one refers to segments, the other refers to features), they can only be in conflict when the relevant feature-to-segment constraints are obeyed.

For instance, if one constraint, Final-[a], references [a], and another constraint, Harmony, references [-ATR] vowels as a class, then they will only be in conflict if [a] is consistently classified as [-ATR]. If the constraint [a] → [-ATR] is always obeyed, then the two constraints Final-[a] and Harmony may impose truly contradictory demands. This is shown in the series of tableaux below, in which [a] → [-ATR] is undominated, but Final-[a] and Harmony have different rankings.

Final-[a]: Assign one violation mark for every word-final vowel which is not [a].

Harmony: Assign one violation mark for every sequence of vowels which have different values for [±ATR].

(6) High ranking for [a] → [-ATR]: Harmony and Final-[a] are in competition[1]

a. Harmony >> Final-[a] gives different final vowel

[+ATR] [-ATR] /b e d a /	[a] → [-ATR]	Harmony	Final-[a]
[+ATR] [+ATR] b e d e			*
[+ATR] [-ATR] b e d a		*!	
[+ATR] [+ATR] b e d a	*!		

1 In the following tableaux, as in all remaining tableaux in this paper, winners will be indicated by shading of the cell instead of by a hand symbol – this is done for reasons of lay-out.

b. Final-[a] >> Harmony gives disruption of harmony to enforce the need to have final [a]

[+ATR] [-ATR] /b e d a /	[a] → [-ATR]	Final-[a]	Harmony
[+ATR] [+ATR] b e d e		*!	
[+ATR] [-ATR] b e d a			*
[+ATR] [+ATR] b e d a	*!		

However, if [a] → [-ATR] may be violated, then Final-[a] and Harmony may both be obeyed in the same candidate. This is shown in the series of tableaux below: because [a] → [-ATR] is low-ranked, the final vowel may be both [a] and [+ATR], and thus, the candidate with this "undesirable" segment/feature combination will win under either ranking of Final-[a] and Harmony.

(7) Low ranking for [a] → [-ATR]: Harmony and Final-[a] are not in competition since they can both be satisfied simultaneously

a. Harmony >> Final-[a] leads to winner with final [a] but "wrong" [+ATR] specification on that vowel

[+ATR] [-ATR] /b e d a /	Harmony	Final-[a]	[a] → [-ATR]
[+ATR] [+ATR] b e d e		*!	
[+ATR] [-ATR] b e d a	*!		
[+ATR] [+ATR] b e d a			*

b. Final-[a] >> Harmony leads to the same result

[+ATR] [-ATR] /b e d a /	Final-[a]	Harmony	[a] → [-ATR]
[+ATR] [+ATR] b e d e		*!	
[+ATR] [-ATR] b e d a	*!		
[+ATR] [+ATR] b e d a			*

Since features and segments are both seen as autonomous labels, neither of which necessarily implies the other (i.e., there can be floating features and featureless segments), there need to be independent Faithfulness constraints for segments and for features, which would protect them from all three basic transformations: deletion, insertion and change. The Max(imality) and Dep(endency) families of constraints, as formulated for segments by McCarthy & Prince (1995) and for features by Lombardi (1995) and others, are sufficient to protect segments and features as independent entities from deletion and insertion in various contexts.

However, since segments are no longer exhaustively defined in terms of features, standard Identity constraints are not sufficient to guarantee lack of change between an input segment and its output correspondent. Identity is normally defined as assigning one violation mark for every feature value which is not shared between an input segment and its output correspondent. Satisfaction of all Identity constraints will entail lack of change between input and output in standard representations. However, in the differential classification model, a mapping such as the following is allowed, in which an input segment value is changed while all its features are retained faithfully from the input:

(8)

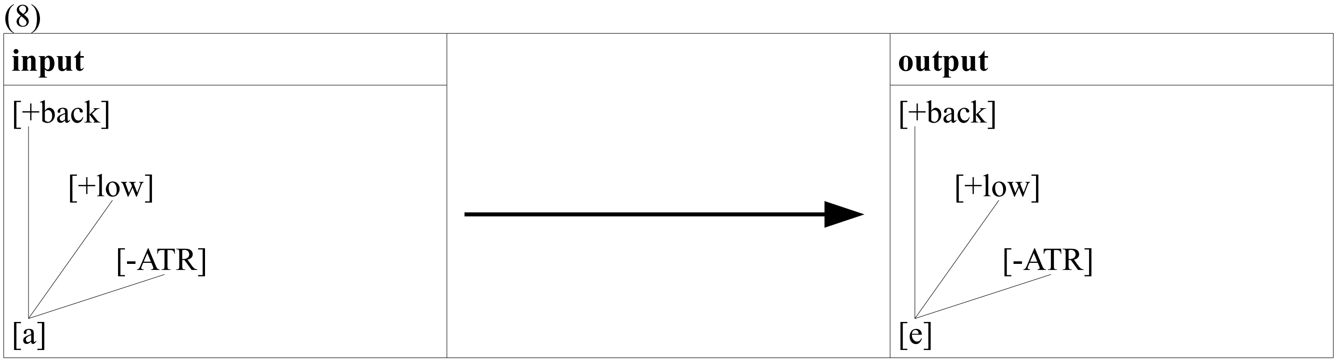

Such a mapping would satisfy all possible Identity constraints, since the output [e] has all the features of input [a]; however, the actual segment value (and with it, the pronunciation) has changed. Of course, a mapping like this could be prevented for winning if both Identity constraints and feature-to-segment constraints were given a high ranking. However, one could imagine a situation in which the connection between, for instance, [e] and [+ATR] must be violable in a certain language (because of opacity, for instance – see the analysis of Yoruba in section 2 of this paper), but there are still independent reasons to prevent segments from randomly changing their value.

Because of such situations, and more broadly, because segments and features are simply two different kinds of representations in the current model, I propose to replace Identity constraints with the family NoChange. Constraints of this family do the same work as Identity constraints do in OT with standard representations: preventing change between corresponding elements in inputs and outputs. However, NoChange constraints look at a pair of input and output elements which correspond to one another, and simply evaluate whether these two elements have the same value. For instance, consider the following definitions:

NoChange([a]): Assign one violation mark for every input segmental value [a] whose output correspondent is not [a].

NoChange([-ATR]): Assign one violation mark for every input [±ATR] feature value [-ATR] whose output correspondent is not [-ATR].

The mapping in (8) above would violate the constraint NoChange([a]), since the segmental value [a] in the input is changed into the segmental value [e]. However, the same mapping would satisfy the constraint NoChange([-ATR]), since input value of [+ATR] has been maintained in the output. Notice that NoChange([-ATR]) does not apply to underlying [+ATR] values.

The table in (9) below gives an overview of mappings which do and do not violate NoChange([-ATR]). The constraint is violated whenever an underlying [-ATR] value is coindexed with an output element which is not [-ATR] (given that coindexation can only occur within a tier, this output element can only be [+ATR]). No violation is incurred in the case of deletion of the feature, change in the underlying segment value (if any) associated with the feature, or a non-canonical segment value (such as [i] with [-ATR]), as can be seen in the first column of the table.

(9)

satisfies NoChange([-ATR])	**violates NoChange([-ATR])**
[-ATR]$_i$ [-ATR]$_i$ │ → │ / a / [a]	[-ATR]$_i$ [+ATR]$_i$ │ → │ / a / [a]
[-ATR]$_i$ [-ATR]$_i$ │ → │ / i / [i]	[-ATR]$_i$ [+ATR]$_i$ → / a / [i]
[-ATR]$_i$ │ → / a / [a]	[-ATR]$_i$ [+ATR]$_i$ → / i / [i]
[-ATR]$_i$ [-ATR]$_i$ │ → │ / a / [i]	
[+ATR]$_i$ [-ATR]$_i$ │ → │ / i / [i]	

Thus, if we restrict our view to just the representational elements [a] and [-ATR], there will be at least the following 6 faithfulness constraints targeting these elements:

Max([a])
Max([-ATR])
Dep([a])
Dep([-ATR])
NoChange([a])
NoChange([-ATR])

In conclusion, the proposal of differential classification requires a slight change in the kinds of constraints that are used in Optimality Theoretical analyses. First, there must be constraints which relate featural and segmental categories to one another (these constraints essentially regulate taxonomic classification of sound events); by their very nature, such constraints are violable, making it possible for non-canonical feature-to-segment associations to surface as outputs of the grammar. Second, regular Markedness constraints may refer to either level of representation used here (featural or segmental), and whether two constraints conflict is determined not only by the configurations targeted by these constraints, but also by which levels the constraints refer to and by the ranking of feature-to-segment association constraints. Lastly, Faithfulness works best in this model if Identity constraints are replaced by NoChange constraints, which do the same work (preventing change between corresponding elements in input and output) as Identity constraints but are formulated to look only within the same tier (is a particular segment value the same between input and output? is a particular [±ATR] value the same between input and output?).

1.3 Phonological opacity in differential classification

Now that the mechanics of the current implementation of the differential classification model have been discussed, I will turn to a discussion of how this model accounts for phonological opacity. One version of the standard definition of phonological opacity (first formulated in Kiparsky 1973) runs as follows[2]:

A phonological rule P of the form A → B / C_D is **opaque** if there are surface structures with the following characteristics:

a. instances of A in the environment C_D. [the rule is not surface-true]
b. instances of B derived by P that occur in environments other than C_D. [the rule is not surface-apparent]
c. instances of B not derived by P that occur in the environment C_D. [the rule is not surface-apparent]

(after McCarthy 2007a:108 (9))

An interaction of rules in which one rule causes the other rule to be opaque (i.e., non-surface-true or non-surface-apparent) is called an opaque interaction, and the phenomenon of sets of phonological data containing an opaque interaction is called phonological opacity.

Definitions of opacity such as the one above refer to phonological rules, which are not components of constraint-based (Optimality Theoretic) analyses. However, it has been observed (see, for instance, McCarthy 1999:356, 362) that situations which would involve a non surface-true or non surface-apparent rule under a rule-based analysis can usually not be derived in standard OT with standard representations: as a rule, only transparent (non-opaque) interactions of generalizations can win in a tableau (see Baković 2011 for classes of opaque interactions for which standard OT does have an account).

One reason why standard OT with standard representations is unable to account for most cases of opacity is because it has no intermediate stages of derivation between underlying and surface form: it is these intermediate stages at which a non-surface-true or non-surface-apparent generalization might be true or apparent. To amend this, extensions to standard OT such as Stratal OT (Kiparsky 2000,

2 Problems with definitions of this kind are discussed in Bakovic (2011).

Bermúdez-Otero 1999) and OT with Candidate Chains (McCarthy 2007b) propose that such intermediate stages are, in fact, available to Optimality Theoretic computation (more will be said about Stratal OT and OT with Candidate Chains in section 4). However, the approach that will be taken in this paper will not involve intermediate stages between input and output; rather, I will take a representational approach.

As opposed to the derivational view of opacity as given in the definition above, opacity can also be seen from a different viewpoint which does not involve rule-based thinking. An alternative definition of opacity, as viewed from a representational rather than derivational perspective, has been suggested by Marc van Oostendorp (p.c.):

Opacity is a situation in which two generalizations are allowed to treat one and the same representational element as having two (seemingly) contradictory statuses.

"Contradictory statuses" may be taken to mean "absence versus presence of a segment", or "contradictory values of the same feature", or "incompatible values of distinct features". For instance, consider the case of Canadian Raising interacting with flapping in certain dialects of North American English (Joos 1942, Chomsky 1964, Bermudez-Otero 2003, 2004, Pater to appear), where the change of /t, d/ to ultra-short (tapped) [d̯] in certain contexts makes the Canadian Raising generalization (diphthongs /aɪ, aʊ/ raise to [əɪ, əʊ] before voiceless consonants) non-surface apparent:

(10)

a. /ɹaɪt/ → [ɹəɪt] "write"	/ɹaɪd/ → [ɹaɪd] "ride"	*(Canadian Raising)*
b. /pæt-ɪŋ/ → [pæd̯ ɪŋ] "patting"	/pæd-ɪŋ/ → [pæd̯ ɪŋ] "padding"	*(flapping)*
c. /ɹaɪt-əɹ/ → [ɹəɪd̯ əɹ] "writer"	/ɹaɪdəɹ/ → [ɹaɪd̯ əɹ] "rider"	*(opaque interaction)*

In this case, it is the element directly following the diphthong which unites contradictory statuses within itself: underlying /t/ which surfaces as [d̯] behaves as [-voice] for the sake of Canadian Raising, but as phonetically voiced [d̯] for the sake of the constraint which bans voiceless coronal stops in certain positions.

When it is a process of deletion or insertion, rather than one of change, that makes another process opaque, then the contradictory properties united into one representational element are presence and absence. For instance, consider the counterfactual scenario in which Canadian Raising were made opaque by a process of word-final deletion of /t/:

(11)

a. /laɪs/ → [ləɪs] "lice"	/laɪz/ → [laɪz] "lies"	
b. /bæt/ → [bæ] "bat"	/bæd/ → [bæd] "bat"	*(t-deletion – hypothetical)*
c. /ɹaɪt/ → [ɹəɪ] "write"	/ɹaɪ/ → [ɹaɪ] "rye"	*(hypothetical opaque interaction)*

In this hypothetical opaque interaction, which is represented by the first item in (...c), underlying /t/ which is deleted is counted as being present for the sake of the constraint which enforces raising of /aɪ/ (because if /t/ were not present, there would be no reason for /aɪ/ to raise), but the same underlying /t/ is counted as absent for the purpose of the constraint which prohibits final [t].

Of course, such contradictory properties could not be united in the same representational element in standard representations. However, differential classification does provide for possibilities to assign contradictory properties to the same sound event, as has been established above. For instance, taking the first Canadian Raising example above, if the feature value [+voice] stands for a set of segments that includes [dl̯], it should be preferably assigned to every instance of of [dl̯], but because the constraint which relates [dl̯] to [+voice] is violable, some constraint rankings will allow [-voice] to be assigned to [dl̯].

I will now sketch a brief analysis of the interaction of Canadian Raising with flapping using the assumptions of differential classification. In particular, I will assume that flapped instances of /t/ are [ɾ] on the segment level, while they are [-voice] on the feature level (as had already been said in the preceding paragraph); in addition, I will assume that the process of flapping refers to the segmental level, while Canadian Raising refers to the featural level. These assumptions, as I will show below, are sufficient to provide a standard OT account of the opaque interaction at hand.

Let us assume that flapping is triggered by the following constraint:

*v[t]v̆ : Assign one violation mark for every instance of the segment [t] which occurs in between two vowels, of which the second is unstressed.

This constraint refers only to the segment [t], not to any of its features (the constraint does refer to features of the surrounding vowels). The formulation of this constraint is simplified: the segment following a flapped segment may also be a consonantal nucleus ([l] or [ɹ]). However, this is not essential to the argument here.

I will assume here that flapping is essentially neutralization of both alveolar stops in English to [d]; I will regard the difference between [d] and [dl̯] as one that is not represented at the segment tier, but rather one that is implemented at a more fine-grained level[3]. However, I assume that the change from [t] to [d] will only involve a change on the segmental tier, and no change on any featural tier. This can be expressed by ranking the constraint NoChange([t]) below *v[t]v̆ , and ranking the blanket constraint Faith(feature), which penalizes any deviation from underlying feature values, above NoChange([t]).

NoChange([t]): Assign one violation mark for any underlying [t] segment which is coindexed with a non-identical output segmental value.

The ranking proposed above (*v[t]v̆ , Faith(feature) >> NoChange([t])) can be justified by the following winner-loser comparisons given the input /pætɪŋ/ "patting":

3 It is possible to represent the difference between [d] and [dl̯] as different segment values, but this would require a slightly more complicated analysis.

(13)

[-voice] /p æ t ɪ ŋ/	*v[t]w̆	Faith(feature)	NoChange([t])
[-voice] p æ d ɪ ŋ			*
[-voice] p æ t ɪ ŋ	*W		L
[+voice] p æ d ɪ ŋ		*W	L

The second candidate in this tableau undergoes no change, thus avoiding the violation of NoChange([t]) present in the winning candidate; however, this candidate violates higher-ranked *v[t]w̆ , and is thus excluded. The third candidate changes both the segmental and the voicing value of [t, -voice], but this variant is excluded by Faith(feature).

Of course, the fact (or rather, assumption) that the winning candidate for this input has a segment [d] which associated with [-voice] means that the constraint which desires for [d] to have the feature value [+voice] must be ranked below both *v[t]w̆ and Faith(feature): otherwise, the current winning candidate would be blocked because it violates this constraint.

[d] → [+voice]: Assign one violation mark for every instance of [d] which is not associated with [+voice].

(13)

[-voice] /p æ t ɪ ŋ/	*v[t]w̆	Faith(feature)	NoChange([t])	[d] → [+voice]
[-voice] p æ d ɪ ŋ			*	*
[-voice] p æ t ɪ ŋ	*W		L	L
[+voice] p æ d ɪ ŋ		*W	L	L

With regard to the process of Canadian Raising, I assume that its triggering context is the feature [-voice]: before [-voice], the diphthongs [aɪ, aʊ] are not allowed.

*[aɪ/aʊ][-voice]: Assign one violation mark for every sequence of the diphthong [aɪ] or [aʊ] and a [-voice] event.

This constraint must be ranked above the constraint which militates against changing these diphthongs: NoChange([aɪ/aʊ]).

NoChange([aɪ/aʊ]): Assign one violation mark whenever an underlying diphthong [aɪ/aʊ] has a non-identical output correspondent.

The following tableau shows the necessity of this ranking (*[aɪ/aʊ][-voice] >> NoChange([aɪ/aʊ])):

(14)

[-voice] /ɹ aɪ t/	*[aɪ/aʊ][-voice]	NoChange([aɪ/aʊ])
[-voice] ɹ əɪ t		*
[-voice] ɹ aɪ t	*W	L

The second candidate violates the higher-ranked markedness constraint *[aɪ/aʊ][-voice], which is why the first candidate wins.

Given these rankings, it can already be seen how the opaque interaction seen in /ɹaɪt-əɹ/ → [ɹəɪdl̯ əɹ] “writer”: the constraint which triggers flapping can only compel a change in the segmental value of underlying [t], but the underlying [-voice] value remains, and it is this value which triggers Canadian Raising, even though the [-voice] value is not reflected in the pronunciation of the word. This is can be seen in the tableau below, where the mapping /ɹaɪt-əɹ/ → [ɹəɪdl̯ əɹ] “writer” is derived without adding any constraints or ranking statements to the ones established above:

(15)

[-vce] /ɹ aɪ t ə ɹ/	*[aɪ/aʊ][-voice]	*v[t]w̆	Faith (feature)	NoChange ([aɪ/aʊ])	NoChange ([t])	[d] → [+voice]
[-vce] ɹ əɪ d ə ɹ				*	*	*
[-vce] ɹ aɪ d ə ɹ	*!				*	*
[+vce] ɹ aɪ d ə ɹ			*!		*	
[+vce] ɹ əɪ d ə ɹ			*!	*	*	
[-vce] ɹ əɪ t ə ɹ		*!		*		
[-vce] ɹ aɪ t ə ɹ	*!	*				

The winning candidate has raising (and thus violates NoChange([aɪ/aʊ])) and flapping (which violates NoChange([t])), but the underlying [-voice] specification of the /t/ remains in place (which violates [d] → [+voice]). All attempts to remove one of this candidate's violations result in violation of higher-ranked constraints.

The second and third candidate in the tableau represent possibilities to have a transparent (bleeding) interaction between Canadian Raising and flapping: the fact that the medial consonant is phonetically voiced prevents Canadian Raising from applying. The second candidate does this by simply not raising the diphthong, while the underlying [-voice] is still maintained on the medial consonant. This leads to a violation of high-ranked *[aɪ/aʊ][-voice]: there is an unraised diphthong before a [-voice] event. The other possibility is changing the underlying [-voice] value on the medial consonant to [+voice], in accordance with [d] → [+voice]; however, this is prevented by high-ranked Faith(feature).

The fourth candidate has raising, but also obeys the feature-to-segment association constraint [d] → [+voice] by assigning [+voice] to the flapped medial consonant. This candidate is also excluded by Faith(feature), but in addition, it has an unnecessary violation of NoChange([aɪ/aʊ]) – the incentive for raising (a following [-voice] event) is not present. Therefore, this candidate is harmonically bounded by the third candidate.

The fifth candidate has raising, but no flapping. This candidate is excluded because it violates the constraint that induces flapping: *v[t]w̆ . Finally, the last and fully faithful candidate has no raising and no flapping, and this candidate is excluded because it violates both *v[t]w̆ and *[aɪ/aʊ][-voice]. In this manner, the candidate which violates both NoChange constraints and the feature-to-segment association constraint (thus, the candidate which has both raising and flapping but an unconventional combination of [d] and [-voice]) is the winning candidate.

Thus, I have shown how opacity can be derived under the assumption of differential classification. Two processes may interact opaquely when one process refers to the featural representation of some sound event, while the other process refers to the segmental representation of the same sound event, and the constraint which enforces the typical feature-segment association for that sound event is low-ranked. In this manner, opacity is a completely expected outcome in certain situations, predicted by the factorial typology given a certain constraint set.

To spell out the potential of differential classification in explaining opaque interactions, I will now present two more detailed case studies of opaque interactions, as observed in Yoruba and Bedouin Arabic, respectively. These case studies constitute sections 2 and 3 of this paper. After this, the paper will be concluded by a section which will compare this account of opacity with other theories, and offer concluding remarks.

2. Case study I: Yoruba

2.1 The data: vowel harmony and vowel assimilation

In this section, I will illustrate how the principle of differential classification accounts for opacity on the basis of a case study of vowel assimilation and vowel harmony in Yoruba. It is well known that Yoruba has Advanced Tongue Root (ATR) vowel harmony (Archangeli & Pulleyblank 1989). However, the generalizations of ATR harmony are systematically violated by the outcome of a process which makes adjacent vowels identical (Pulleyblank 1988). I will show that differential classification provides an elegant account of this opaque interaction.

Yoruba (Bamgbose 1966, Pulleyblank 1988, Archangeli & Pulleyblank 1989, Nevins 2004) exhibits vowel harmony on the dimension of Advanced Tongue Root (ATR): within a word, certain combinations of vowels in subsequent syllables are banned if they differ in their [±ATR] value. Yoruba has the following vowel inventory:

(16)

i	ĩ		u
e	ẽ		o
ɛ			ɔ
		a	ã

Following Pulleyblank (1988) and Archangeli & Pulleyblank (1989), I will restrict the present discussion to the oral vowels [i, u, e, o, ɛ, ɔ, a]; the nasalized vowels [ĩ, ẽ, ã] appear to behave like their oral counterparts [i, e, a], so that, for the sake of simplicity of exposition, nasal vowels will be treated as “variants” of their oral counterparts.

The high and high-mid vowels [i, u, e, o] (shown on a shaded background above) are considered to be [+ATR] in surface representations by Archangeli & Pulleyblank (1989), while the other vowels [ɛ, ɔ, a] are considered to be [-ATR]. The particular sequences of vowels (in adjacent syllables within the same word) that are banned are:

(17)

a. *e ... a	b. *o ... a
c. *e ... ɛ	d. *o ... ɔ
e. *e ... ɔ	f. *o ... ɛ
g. *ɛ ... e	h. *ɔ ... o
i. *ɛ ... e	j. *ɔ ... e

Thus, a mid [+ATR] vowel may not precede any [-ATR] vowel [a, ɛ, ɔ] (compare (16a-f)), while it may also not follow a mid [-ATR] vowel (as in (17e,f)). These bans are only valid of vowels in adjacent syllables.

High vowels may co-occur with any [ATR] value in an adjacent syllable. High vowels also block harmony, in the sense that vowels left and right from a high vowel need not correspond in their ATR value, as can be seen in (18b). This contrasts with low vowels, which may occur *before* any [ATR] value, but are restricted to occurring only after [-ATR] vowels – as shown in (18c).

(18)
a. ɛbi "guilt"
   ɔkĩ "[r]egret"
b. elubɔ "yam flour"[4][5]
   odidɛ "gray parrot"
c. ɛrakpo "type of plant"[6]

In words with more than one morpheme, it can be seen that [ATR] harmony spreads from roots to prefixes. This may be analyzed as left-to-right directionality:

(19)
/o-ku/ → [oku] "corpse of a person"
/o-dɛ/ → [ɔdɛ] "hunter"
/e-ro/ → [ero] "a thought"
/e-rɔ/ → [ɛrɔ] "machine"

I will not attempt to account for the left-to-right directionality of harmony. However, it must be kept in mind that ATR harmony has both a phonotactic (shape of individual morphemes) and a dynamic (root-to-prefix spread) side.

Beside vowel harmony, Yoruba also has an (optional) process of full vowel assimilation[7]: whenever two vowels are absolutely adjacent, one of the vowels takes on the same quality as the other vowel. I will restrict the current analysis to adjacent vowels which occur at word boundaries (Yoruba normally allows only CV syllables word-internally, so that adjacent vowels are rare inside a word: they may arise through consonant deletion, or reduplication).

(20)
/ara ilu/ → [ara alu] "townsman"

Full vowel assimilation is regressive (i.e., right-to-left), except when the second vowel is [i], in which case it is progressive (left-to-right). In fact, whenever (standard) regressive assimilation occurs, the second vowel is never a high vowel: since words starting in /u/ do not occur in Yoruba, the second vowel will never be underlyingly /u/. The examples below demonstrate this.

(21)
a. /awo ejo/ → [aw**e** ejo] "color of a snake"
   /ara oke/ → [ar**o** oke] "Northern Yoruba"

4 Tone will be omitted in all Yoruba examples cited here.
5 All examples cited are from Archangeli & Pulleyblank (1989) and Pulleyblank (1988).
6 Baković (2003:15 (20a.ii))
7 Pulleyblank (1988) describes assimilation as a set of two distinct rules: Regressive Feature Spread and Progressive Feature Spread. However, OT makes it possible to analyze the directionality of vowel assimilation without splitting the process into two different subprocesses.

/ara *uke/

b. /ara ilu/ → [ara **a**lu] “townsman”

/ɛru igi/ → [ɛru **u**gi] “bundle of wood”

To restrict the scope of the current analysis, directionality of vowel assimilation will not be given a formal analysis. However, these facts of directionality support the idea that high vowels in Yoruba lack certain features (Pulleyblank 1988; essentially, the lack of features on [i] prevents its vowel quality from being spread), and this idea will be part of my analysis of ATR harmony.

The process of full vowel assimilation freely creates exceptions to ATR harmony, which is opaque behavior. Vowel assimilation causes a vowel quality from one word to be copied into another word, and this vowel quality copying takes places regardless of whether this vowel quality is appropriate with respect ATR harmony: a [+ATR] vowel might be copied into a position where only a [-ATR] vowel would be appropriate, or vice versa. Yoruba does not avoid such situations: disharmony created by vowel assimilation is allowed, even though disharmony is avoided otherwise. The examples below demonstrate how full vowel assimilation overrides ATR harmony.

(22)

/ɛba odo/ → [**ɛbo** odo] “near the river” (*[ɛ ... o] is normally not allowed)

/ile iʃɛ/ → [ile **eʃɛ**] “office” (*[e ... ɛ] is normally not allowed)

/owo ɔmɔ/ → [**owɔ** ɔmɔ] “child's money” (*[o ... ɔ] is normally not allowed)

This interaction between ATR harmony and vowel assimilation may be characterized as a non-surface-true interaction: whenever vowel assimilation creates a disharmonic sequence, there is an opportunity for dynamic right-to-left vowel harmony to apply (which would be a feeding interaction between vowel assimilation and vowel harmony), but this does not happen.

(23) *hypothetical feeding interaction:*

/ɛba odo/ → *[**e**bo odo] (vowel assimilation plus leftward harmony)

However, if ATR harmony is considered in its phonotactic aspect, that the interaction between assimilation and harmony can also be seen as a non-surface-apparent interaction. Whenever vowel assimilation threatens to create word-internal sequences which conflict with ATR harmony, phonotactics could block the process of vowel assimilation. This also does not occur.

(24) *hypothetical blocking interaction:*

/awo ejo/ → [awe ejo], *but*

/ɛba odo/ → *[ɛba odo] *(no vowel assimilation because it would otherwise lead to disharmony)*

A traditional, rule-based approach to this interaction would be to order the ATR harmony rule before the vowel assimilation rule. This would prevent the ATR harmony rule from changing the outcome of the vowel assimilation rule:

(25)

input	/Eba odo/
1. harmony rule	ɛba odo
2. vowel assimilation rule	ɛbo odo
~~harmony rule~~	- (~~ebo odo~~)
output	[ɛbo odo]

By contrast, the account I propose here for the opaque interaction between vowel harmony and vowel assimilation will be built on the assumption that vowel harmony operates at a more abstract level than vowel assimilation. To be precise, I assume that the process of ATR harmony refers to the feature [±ATR] and not to segmental labels, while vowel assimilation only involves changing values at the segmental level, and does not refer to the feature [±ATR] at all.

The assumption that harmony is represented in the grammar using features, while assimilation is represented using atomic segments can actually be justified by the nature of the data: this assumption makes it possible to analyze the patterns with a fewer number of constraints. Given that ATR harmony is based on the (in)compatibility of classes of segments (for instance, the incompatibility of [ɛ, ɔ] with [e, o]), it can be represented with fewer and simpler constraints if reference is made to features that refer to ATR and RTR vowels as classes: the analysis below has two constraints which trigger ATR harmony. The alternative, representing each disallowed combination of vowels as a separate constraint on co-occurrence of segment labels, would require at least 10 different triggering constraints to be formulated[8].

Similarly, the requirement that two adjacent vowels be fully identical, which fuels the process of vowel assimilation, can be expressed by one single constraint if this constraint refers to holistic segment identity: this constraint may simply state that two adjacent values on the segment tier must be identical whenever they are both classified as vowels. On the other hand, if the constraints triggering assimilation only referred to feature values, then there should be a separate constraint requiring identity of two adjacent values on a particular feature dimension (requirement that they have the same backness, the same ATR value, the same height, etc.), which would lead to more than one constraint.

The notion that it is desirable to represent a pattern with fewer and simpler constraints is intuitive, but has no formal backing. One way to bring this intuition into the realm of the formal is by resurrecting the evaluation metric for phonological grammars as proposed by Chomsky & Halle (1968), which could roughly be stated as follows: whenever a certain phonological phenomenon can be grammatically represented in more than one way, the grammar which employs the fewest number of symbols is chosen. This principle will generally lead to a maximally concise and general statement of every phonological pattern in the grammar. In section 4 of this paper, I will sketch an reformulation of this evaluation metric in terms of constraint induction, which will provide a principled reason to choose the grammatical representation of Yoruba harmony and assimilation as assumed above over other potential ways of representing it.

8 This, of course, presumes that constraint definitions would normally not contain disjunctions such as “Assign one violation mark whenever two consecutive vowels are [e] and [ɛ] OR [o] and [ɔ] OR [e] and [a] OR ...”. The inclusion of disjunctions into constraint definitions seems generally undesirable given the goal of OT to represent patterns by using interaction of individual constraints.

If it is true that the constraints which drive harmony and those that drive assimilation refers to different taxonomical levels, then, according to the general schema outlined in section 1.4, the two processes may interact opaquely, provided that the constraints relating the featural and segmental values that are involved to one another are sufficiently low-ranked. In the remainder of this section, I will show how such an account derives the Yoruba data.

2.2 A differential classification account

The vowel system of Yoruba, as shown above, must be classified along several featural dimensions (Archangeli & Pulleyblank 1989). The dimension relevant here is that of Advanced Tongue Root (ATR). Following Pulleyblank (1988), I assume that whether and how vowels are classified on this featural dimension is determined by their height. The high vowels [i] and [u] are assumed to never carry a [±ATR] specification, which means that the following constraint is undominated:

*[i/u, ±ATR]:

> Assign one violation mark for every segmental value [i] or [u] which is associated with a value of the feature [±ATR].

I also assume that [e, o] are standardly associated with [+ATR], and [ɛ, ɔ] and [a] are standardly associated with [-ATR]. This is expressed by the following constraints:

{e, o} → [+ATR]:

> Assign one violation mark for every segmental value [e] or [o] which is not associated with [+ATR].

{ɛ, ɔ} → [-ATR]:

> Assign one violation mark for every segmental value from the set {ɛ, ɔ} which is not associated with [-ATR].

[a] → [-ATR]:

> Assign one violation mark for every segmental value [a] which is not associated with [-ATR].

As has been said above, I assume that vowel assimilation involves a ban on adjacent non-identical vowel segments, rather than a ban on adjacent vowels whose features are not identical. The constraint which triggers vowel assimilation should refer to the segmental tier, as well as the [±vocalic] tier. Given this, the constraint which triggers vowel assimilation can be formulated as follows:

AssimV:

> Assign one violation mark for every sequence of two adjacent [+vocalic] specifications which are associated with distinct values on the segmental tier.

This constraint only punishes configurations such as the one in (26a), while configurations such as those in (26b) do not violate AssimV.

(26)

a. Banned	b. Allowed
[+voc][+voc] [a] [i]	[+voc][+voc] [+voc][-voc] [a] [a] [a] [j]

Because the constraint AssimV is enforced by changing one of two adjacent vowel qualities, AssimV compels violations of, and therefore dominates, the faithfulness constraint NoChange(segment):

NoChange(segment):

> Assign one violation mark for every input specification on the segment tier whose output correspondent is not identical to it.

In the tableau below, this ranking is shown on the basis of the example /ara ilu/ → [ara alu]:

(27)

[+v] [+v] [+v] [+v] / a r a i l u /	AssimV	NoChange(segment)
[+v] [+v] [+v] [+v] a r a a l u		*
[+v] [+v] [+v] [+v] a r a i l u	*W	L

The tableau below shows how the ranking of this constraint above NoChange(segment) prevents AssimV from being satisfied by configurations with lacking [+vocalic] specifications on vowels. The input used here is purely hypothetical (does not necessarily correspond to a pair of existing Yoruba words), and lacks a number of [+vocalic] specifications:

(28)

[+v] [+v] / a gb a i r u/	{a,e,ɛ,o,ɔ,i,u} → [+vocalic]	AssimV	NoChange(segment)
[+v] [+v] [+v] [+v] a gb a a r u			*
[+v] [+v] [+v] [+v] a gb a i r u		*W	L
[+v] [+v] a gb a i r u	**W		L

The second candidate in this tableau has [+vocalic] values on each vowel (I will ignore the fact that this violates faithfulness constraints on values of the feature [±vocalic]), but retains underlying segment qualities, which leads to a violation of Assim(V). The third candidate lacks a [+vocalic] value on the first of two adjacent vowels, making it impossible for this [a i] sequence to violate AssimV; however, this candidate is ruled out by the high-ranked constraint against vowel segments without a [+vocalic] specification.

One issue that remains is the directionality of vowel assimilation. However, I will defer its discussion until the end of this section. Before that, the analysis of ATR harmony, the process with which vowel assimilation makes opaque, will be discussed.

As was shown in the previous subsection, ATR harmony in Yoruba depends on the dimensions of ATR and vowel height. The exact pairs of vowels which are not allowed in adjacent syllables (see (29) below) can be summarized as follows (as has already been said in the previous section):

1) non-high [+ATR] vowels [e, o] may not precede [-ATR] vowels [ɛ, ɔ, a];
2) mid [-ATR] vowels [ɛ, ɔ] may not precede non-high [+ATR] vowels [e, o].

(29) *Disallowed (disharmonic) vowel combinations in Yoruba* (=(17) above)

a. *e ... a	b. *o ... a
c. *e ... ɛ	d. *o ... ɔ
e. *e ... ɔ	f. *o ... ɛ
g. *ɛ ... e	h. *ɔ ... o
i. *ɛ ... e	j. *ɔ ... e

These generalizations are obtained through the interaction of two simple harmony constraints with various feature-to-vowel assignment constraints. The harmony constraints can be formulated in the following way:

*[-ATR][+ATR]:

> Assign one violation mark for a word-internal pair of syllables, of which the first syllable's head is associated with[-ATR] and the second syllable's head is associated with [+ATR].

*[+ATR][-ATR]:

> Assign one violation mark for a word-internal pair of syllables, of which the first syllable's head is associated with [+ATR] and the second syllable's head is associated with [-ATR].

Because high vowels are banned from having a [±ATR] value as per the undominated constraint *[i/u, ±ATR] introduced at the beginning of this section, high vowels may occur next to both [+ATR] and [-ATR] vowels. The input /odidɛ/ is mapped onto [odidɛ], because having a [±ATR] value on the middle vowel [i] is excluded – otherwise the output could have been [ɔdidɛ] – see the second candidate in the tableau, which is the only candidate where [i] has a [±ATR] value and the harmony constraints are obeyed.

(30)

[+atr] [-atr] / o d i d ɛ /	*[i/u, ±ATR]	*[-ATR][+ATR]	*[+ATR][-ATR]
[+atr] [-atr] o d i d ɛ			
[-atr] [-atr] [-atr] ɔ d i d ɛ	*!		
[+atr] [-atr] [-atr] o d i d ɛ	*!		*
[+atr] [+atr] [-atr] o d i d ɛ	*!		*

The fact that [a] can only be preceded by [-ATR] or high vowels, but may be followed by any vowel, can be accounted by making the requirements that underlying [a] be preserved and have the feature [-ATR] more important than vowel harmony. The ranking [a] → [-ATR] >> *[-ATR][+ATR] makes it possible for candidates which contain [a] and a following [+ATR] vowel to disagree, which is necessary to obtain the grammatical form of words such as [ɛrakpo] “type of plant”.

(31)

[-atr] [-atr] [+atr] / ɛ r a kp o /	[a] → [-ATR]	*[-ATR][+ATR]
[-atr] [-atr] [+atr] ɛ r a kp o		*
[+atr] [+atr] [+atr] e r a kp o	*W	L

Of course, for the desired candidate to win in this comparison, some faithfulness constraint must protect underlying [a] from changing to some other vowel quality which may be [+ATR]. Let us assume this faithfulness constraint is NoChange([a]):

NoChange([a]):

> Assign one violation mark for every segmental value [a] whose output correspondent is not identical to it.

This faithfulness constraint (along with [a] → [-ATR]) must be ranked above *[-ATR][+ATR], so that underlying [a] will not surface as some other vowel for the sake of vowel harmony:

(32)

[-atr] [-atr] [+atr] / ɛ r a kp o /	NoChange([a])	[a] → [-ATR]	*[-ATR][+ATR]
[-atr] [-atr] [+atr] ɛ r a kp o			*
[+atr] [+atr] [+atr] e r e kp o	*W		L
[+atr] [+atr] [+atr] e r a kp o		*W	L

To prevent the last vowel ([o]) in this word from being changed to a [-ATR] vowel – which would prevent violation of *[-ATR][+ATR] – faithfulness to [+ATR] must also be ranked above *[-ATR][+ATR]. The relevant faithfulness constraint is NoChange([+ATR]):

NoChange([+ATR]:

> Assign one violation mark for every input [+ATR] value whose output value is not identical to it.

(33)

[-atr] [-atr] [+atr] / ɛ r a kp o /	NoChange([+ATR])	*[-ATR][+ATR]
[-atr] [-atr] [+atr] ɛ r a kp o		*
[-atr] [-atr] [-atr] ɛ r a kp ɔ	*W	L

However, NoChange([+ATR]) must be ranked *below* the other constraint which enforces harmony: *[+ATR][-ATR]. This is because [a] requires being preceded by a [-ATR] vowel:

(34)

[+atr] [-atr] / o j a /	*[+ATR][-ATR]	NoChange([+ATR])
[-atr] [-atr] ɔ j a		*
[+atr] [-atr] o j a	*!	

By contrast, the NoChange constraint referring to [-ATR] must be ranked below both harmony constraints, because both harmony constraints can compel a change from a [-ATR] to a [+ATR] vowel. The general constraint against changing segmental values – NoChange(segment) – must also be dominated by the harmony constraints, since harmony leads to change of segmental values (instead of “covert harmony”, where vowels agree in ATR specification but their pronunciation does not reflect this – which is shown in the second, losing candidate in (35) below).

NoChange(segment):

Assign one violation for any input value on the segmental value whose corresponding output value is not identical to it.

(35)

[-atr] [+atr] / ɛ p o /	*[-ATR][+ATR]	NoChange([-ATR])	NoChange(segment)
[+atr] [+atr] e p o		*	*
[-atr] [+atr] ɛ p o	*W	L	L

At the same time, the constraints which assign [±ATR] values to mid vowels should be ranked above NoChange(segment), so that harmony constraints cannot be satisfied by changing only [±ATR] but not segmental values. In the tableaux below, it is demonstrated that ranking harmony constraints and constraints which assign [±ATR] above NoChange(segment) leads to candidates with disharmony expressed on the segmental level being ruled out.

(36)

[+atr] [-atr] / e d ɔ /	{e, o} → [+ATR]	*[+ATR][-ATR]	NoChange(segment)
[-atr] [-atr] ɛ d ɔ			*
[-atr] [-atr] e d ɔ	*W		L
[+atr] [-atr] e d ɔ		*W	L

(37)

[-atr] [+atr] / ɛ d o /	{ɛ, ɔ} → [-ATR]	*[+ATR][-ATR]	NoChange(segment)
[+atr] [+atr] e d o			*
[+atr] [+atr] ɛ d o	*W		L
[+atr] [-atr] ɛ d o		*W	L

As a final point, it must be mentioned that, if it is assumed that a word like [ɛrakpo] has an underlying form like /erakpo/, then it can be shown that NoChange([a]) and [a] → [-ATR] must be ranked above

NoChange([+ATR]), because it is these two constraints pertaining to [a] which force the first vowel of this word to change from [+ATR] to [-ATR]:

(38)

[+atr] [-atr] [+atr] / e r a kp o /	NoChange([a])	[a] → [-ATR]	NoChange([+ATR])
[-atr] [-atr] [+atr] ɛ r a kp o			*
[+atr] [+atr] [+atr] e r e kp o	*W		L
[+atr] [+atr] [+atr] e r a kp o		*W	L

The constraint hierarchy thus established for ATR harmony is represented in the Hasse diagram below:

(39)

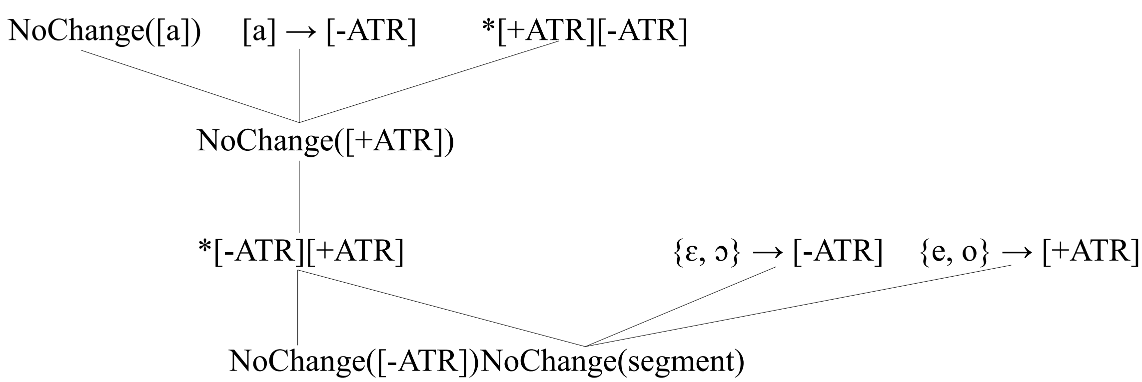

The summary tableaux will show how this ranking derives the facts of Yoruba ATR vowel harmony. The mappings that will be considered are /o-dɛ/ → [ɔdɛ] “hunter”, /ɛ-ro/[9] → [e-ro] “a thought”, /odidɛ/ → [odidɛ] “gray parrot” and /erakpo/[10] → [ɛrakpo] “type of plant”.

9 There is no evidence to assume that the underlying form of the first vowel of this word is [ɛ] with a [-ATR] feature. This underlying form is chosen only for reasons of exposition.

10 There is no evidence to assume that the first vowel of this word is underlyingly [e] and [+ATR], but I will assume this input for expository reasons.

(40)

+atr -atr / o d ɛ/	*[+ATR][-ATR]	NoChange([a])	[a] → [-ATR]	NoChange([+ATR])	*[-ATR][+ATR]	{ɛ, ɔ} → [-ATR]	{e, o} → [+ATR]	NoChange([-ATR])	NoChange(segment)
-atr -atr ɔ d ɛ				*					*
+atr -atr o d ɛ	*!			*					
+atr -atr ɔ d ɛ	*!					*			*
-atr -atr o d ɛ				*			*!		

In this tableau, the winning candidate has the segmental value and [±ATR] value of its first vowel changed to make the entire word a [-ATR] word. The fully faithful (second) candidate has a disharmonic sequence of ATR values, and is excluded by its violation of high-ranked *[+ATR][-ATR]. The third candidate in this tableau, which changes only the segmental value of the first vowel, but not its [±ATR] value, still violates the harmony-inducing constraint *[+ATR][-ATR], and therefore is excluded. Finally, when only the [±ATR] value of the first vowel is changed, but not its segmental value, this candidate is ruled out because it violates {e, o} → [+ATR], even though this obviates a violation of lower-ranked NoChange(segment).

Candidates in which ATR harmony is progressive rather than regressive were not included in this tableau, because the current account explicitly does not explain the direction of vowel harmony, and Yoruba vowel harmony is always regressive (leftward) rather than progressive (rightward). Some candidates with rightward harmony could have been more harmonic than the winner in this tableau (because they would have avoided violation of NoChange([+ATR)]), but I assume that some other constraint(s) would have ruled such candidates out simply because they progressive assimilation.

(41)

-atr +atr / ɛ r o/	*[+ATR][-ATR]	NoChange([a])	[a] → [-ATR]	NoChange([+ATR])	*[-ATR][+ATR]	{ɛ, ɔ} → [-ATR]	{e, o} → [+ATR]	NoChange([-ATR])	NoChange(segment)
+atr +atr e r o								*	*
-atr +atr ɛ r o					*!				
-atr +atr e r o					*!		*		*
+atr +atr ɛ r o						*!		*	

This tableau is structured like the previous one: the second candidate is fully faithful and is excluded by a harmony-inducing constraint (in this case, *[-ATR][+ATR]), while the third candidate only has a change in the segmental value of the first vowel, which still leads to a violation of *[-ATR][+ATR]. Finally, the fourth candidate only has a change in [±ATR] value in the first vowel, and this satisfies the harmony constraint, but this candidate violates the constraint which enforces the classification of [ɛ] as [-ATR], and this constraint is ranked higher than both NoChange constraints violated by the winner.

(42)

+atr -atr / o d i d ɛ/	*[+ATR][-ATR]	NoChange([a])	[a] → [-ATR]	*[i/u, ±ATR]	NoChange([+ATR])	*[-ATR][+ATR]	{ɛ, ɔ} → [-ATR]	{e, o} → [+ATR]	NoChange([-ATR])	NoChange(segment)
+atr -atr o d i d ɛ										
+ar+ar-atr o d i d ɛ	*!			*						
-atr-atr-atr ɔ d i d ɛ				*	*!					*
+atr -atr ɔ d i d ɛ							*!			*

The undominated constraint *[i/u, ±ATR] is included in this tableau because of the presence of [i] in this example. The winning (and fully faithful) candidate has the vowel [i] without a [±ATR] feature on it, flanked by two vowels with contradictory [±ATR] values. However, this winning candidate does not violate any of the constraints in this analysis. Crucially, the contradictory [±ATR] values on the two outer values of this word do not violate either of the harmony constraints, which are defined over adjacent syllables.

Inserting a [±ATR] value on the middle [i], as in the second candidate, actually does trigger a violation of the high-ranked harmony constraint *[+ATR][-ATR]. Avoiding this violation by making the entire word a [-ATR] word, as in the third candidate, results in an unnecessary violation of NoChange (because the fully faithful candidate violates no relevant Markedness constraint). Finally, changing the vocalic identity of the first vowel only (to avoid a violation of NoChange([+ATR])) results in needless violation of {ɛ, ɔ} → [-ATR].

(43)

+atr -atr +atr / e r a kp o/	*[+ATR][-ATR]	NoChange([a])	[a] → [-ATR]	NoChange([+ATR])	*[-ATR][+ATR]	{ɛ, ɔ} → [-ATR]	{e, o} → [+ATR]	NoChange([-ATR])	NoChange(segment)
-atr -atr +atr ɛ r a kp o				*	*				*
+atr -atr +atr e r a kp o	*!				*				
+atr +atr +atr e r a kp o			*!					*	
+atr +atr +atr e r e kp o		*!						*	*
-atr -atr +atr e r a kp o				*	*		*!		

The second, fully faithful, candidate in this tableau is ruled out because its first two syllables violate the highest-ranked harmony constraint *[+ATR][-ATR]. In the third candidate, this problem is tentatively solved by classifying [a] as a [+ATR] vowel, but this violates the equally high-ranked constraint [a] → [-ATR]. In the fourth candidate, both preceding problems (disharmony and feature assignment to [a]) are avoided by changing underlying [a] to [e], which must be a [+ATR] vowel. However, this is excluded by the high ranking of NoChange([a]). Thus, it is the first vowel, not the second vowel of the word which must change. The last candidate of this tableau attempts to change only the feature value of that vowel, but this leads to a mismatch between segmental and featural qualities. This mismatch is punished by a higher-ranked constraint ({e, o} → [+ATR]) than changing the vowel quality of the first vowel (NoChange(segment)), so that the resulting winning candidate has change of the underlying segmental quality of the first vowel, rather than simply a mismatch between the segmental and [±ATR] values of that vowel.

Now that an account of the process of ATR harmony has been developed, I will turn towards the opaque interaction between ATR harmony and vowel assimilation. Let us assume that this opaque interaction arises because mismatches between vowels and their canonical ATR values are allowed. Thus, a vowel like [o] may be classified as [-ATR] in some cases, while a vowel like [ɔ] may sometimes be classified as [+ATR]. In this manner, we can assume that the outputs of opaque mappings such as /ɛba odo/ → [ɛbo odo] will not actually violate the harmony constraints established above: if the segment [o] in the first word [ɛbo] is allowed to have the value [-ATR], despite [o] normally having the value [+ATR], then vowel harmony will be maintained at the featural level, while vowel assimilation is true of the segmental level:

(44)
[-atr] **[-atr]** [+atr] [+atr]
ɛ b o o d o

The question, of course, is how to allow such exceptional mappings just when necessary. To ensure that ATR "mismatches" are allowed in mappings such as /ɛba odo/ → [ɛbo odo], in which vowel matching at word edges is compelled by AssimV, the constraints which regulate assignment of the feature [±ATR] must be ranked below AssimV. This is shown in the tableaux below, of which the first shows how [o] is forced onto the [-ATR] feature of underlying /a/, while the second tableau shows how [ɔ] is forced onto the [+ATR] feature of underlying /o/[11]:

(45)

[-atr] [-atr] [+atr] [+atr] / ɛ b a o d o /	AssimV	{e, o} → [+ATR]
[-atr] [-atr] [+atr] [+atr] ɛ b o o d o		*
[-atr] [-atr] [+atr] [+atr] ɛ b a o d o	*W	L

(46)

[+atr] [+atr] [-atr] [-atr] / o w o ɔ m ɔ /	AssimV	{ɛ, ɔ} → [-ATR]
[+atr] [+atr] [-atr] [-atr] o w ɔ ɔ m ɔ		*
[+atr] [+atr] [-atr] [-atr] o w o ɔ m ɔ	*W	L

Thus, the ranking AssimV >> {e, o} → [+ATR], {ɛ, ɔ} → [-ATR] prevents a blocking relationship between harmony and vowel assimilation (i.e. harmony blocking the application of vowel assimilation), which is one of the two possible transparent interactions between harmony and assimilation discussed in the preceding subsection. The other possible transparent interaction is one of feeding: vowel assimilation creates potential inputs for dynamic regressive ATR harmony, which would lead to potentially changing more than one vowel in the first of two words – for instance, /ɛba odo/ would sound as *[ebo odo], and /owo ɔmɔ/, as *[ɔwɔ ɔmɔ].

To prevent this transparent interaction, the faithfulness constraints NoChange([-ATR]) and NoChange([+ATR]) must also be ranked above the two constraints which regulate assignment of the feature [±ATR]. The mechanism behind this is the fact that there are two possible ways of satisfying ATR harmony when regressive vowel assimilation results in a [+ATR] vowel in a [-ATR] word (or the other way around): allowing an "incorrect" ATR value on the newly created vowel in the first word ([ɛbo[-ATR] odo]), which violates segment-to-[±ATR] constraints, or changing underlying [±ATR] values to accomodate the newly created vowel ([ebo[+ATR] odo]), which violates NoChange([-ATR]). Because the former, opaque, solution is chosen over the latter, transparent, solution in Yoruba, NoChange([±ATR]) must outrank segment-to-[±ATR] constraints.

11 The sources consulted do not actually list examples of the shape /...V aCV/ → [...a aCV], where assimilation of a vowel to [a] leads to opacity of vowel harmony. This is why the status of [a] as potentially having a [+ATR] specification through vowel assimilation is ignored in this discussion.

(47)

[-atr] [-atr] [+atr] [+atr] / ɛ b a o d o /	NoChange([-ATR])	{e, o} → [+ATR]
[-atr] [-atr] [+atr] [+atr] ɛ b o o d o		*
[+atr] [+atr] [+atr] [+atr] e b o o d o	**W	L

(48)

[+atr] [+atr] [-atr] [-atr] / o w o ɔ m ɔ /	NoChange([+ATR])	{ɛ, ɔ} → [-ATR]
[+atr] [+atr] [-atr] [-atr] o w ɔ ɔ m ɔ		*
[-atr] [-atr] [-atr] [-atr] ɔ w ɔ ɔ m ɔ	*W	L

The ranking which results from this is the following (not counting the undominated constraint *[i/u, ±ATR]):

(49)

NoChange([a])   [a] → [-ATR]   *[+ATR][-ATR]

NoChange([+ATR])

*[-ATR][+ATR]

NoChange([-ATR])   AssimV

{e, o} → [+ATR]   {ɛ, ɔ} → [-ATR]

NoChange(segment)

The summary tableaux below show how this ranking derives the opaque interaction between ATR harmony and vowel assimilation. I will cover the cases /ɛba odo/ → [ɛbo odo] “near the river”, /owo ɔmɔ/ → [owɔ ɔmɔ] “child's money” and the constructed case /ɛbe odo/ → [ebo odo], in which vowel assimilation creates no exception to ATR harmony because the underlying form of the first word /ɛbe/

would normally always surface as a [+ATR] word [ebe][12].

(50)

-a -a +a +a /ɛ ba o do/	NoChange([a])	[a] → [-ATR]	*[+ATR][-ATR]	NoChange([+ATR])	*[-ATR][+ATR]	NoChange([-ATR])	AssimV	{e, o} → [+ATR]	{ɛ, ɔ} → [-ATR]	NoChange(segment)
-a -a +a +a ɛ bo o do								*		*
-a -a +a +a ɛ ba o do							*!			
+a+a+a +a e bo o do						*!*				**
-a+a+a +a ɛ bo o do					*!	*				*

This tableau compares the actually observed, opaque, mapping /ɛba odo/ → [ɛbo odo] to the two main transparent countercandidates: the second candidate, in which vowel assimilation is blocked because it would otherwise violate the segment-level reflection of ATR harmony, and the third candidate, in which vowel assimilation feeds ATR harmony. The second candidate is ruled out because AssimV dominates both constraints violated by the winning candidate: NoChange in reference to segment values, and the requirement that [o] have the feature value [+ATR]. The third candidate is ruled out because it requires excessive violations of NoChange in reference to [-ATR] vowels.
Finally, the fourth candidate represents an alternative parse of the observed surface form of this mapping: one in which vowel assimilation has led to change of the [±ATR] value of the second vowel of the first word, but only this particular [±ATR] value. This candidate is ruled out because it actually does violate ATR harmony, as opposed to the winning candidate.

(51)

+a +a -a -a /o wo ɔmɔ/	NoChange([a])	[a] → [-ATR]	*[+ATR][-ATR]	NoChange([+ATR])	*[-ATR][+ATR]	NoChange([-ATR])	AssimV	{e, o} → [+ATR]	{ɛ, ɔ} → [-ATR]	NoChange(segment)
+a +a -a -a o wɔ ɔmɔ									*	*
+a +a -a -a o wo ɔmɔ							*!			
-a -a -a -a ɔ wɔ ɔmɔ				*!*						**
+a -a -a -a o wɔ ɔmɔ			*!	*						*

This tableau is organized in the same way as the one above, except that this mapping involves copying

12 In this series of tableaux, values of [±ATR] will be represented as "+a" and "-a".

a [-ATR] vowel into a [+ATR] position, rather than the opposite (which was the case in the tableau above). However, the grammatical mechanism is the same. The candidate where surface realization of ATR harmony blocks vowel assimilation (which is the second candidate) is ruled out by its violation of AssimV. The candidate with a feeding relation between vowel assimilation and harmony (the third candidate) is ruled out because it excessively violates NoChange([+ATR]). Finally, the fourth candidate violates one of the harmony constraints (*[+ATR][-ATR]) and is excluded for that reason.

(52)

-a +a +a +a /ɛ be o do/	NoChange([a])	[a] → [-ATR]	*[+ATR][-ATR]	NoChange([+ATR])	*[-ATR][+ATR]	NoChange([-ATR])	AssimV	{e, o} → [+ATR]	{ɛ, ɔ} → [-ATR]	NoChange(segment)
+a+a+a +a e bo o do						*				**
-a +a +a +a ɛ be o do					*!		*			
-a +a +a +a ɛ bo o do					*!			*		*
+a+a+a +a e be o do						*	*!			*

The first word in this (constructed) sequence is underlyingly disharmonic. Therefore, the fully faithful (second) candidate is excluded by its violation of high-ranked *[-ATR][+ATR]. The third candidate offers some improvement over the second candidate by assimilating the last vowel of the first word to the first vowel of the second word, thus avoiding violation of AssimV; however, this candidate still violates *[-ATR][+ATR]. Finally, the fourth candidate has the first vowel of the first word harmonized with the second vowel of that word; however, because it has no vowel assimilation, this candidate is excluded by its violation of AssimV. Thus, the remaining and optimal option is to have both ATR harmony and vowel assimilation apply, as in the winning candidate. This mapping shows that the application of vowel assimilation alone is not a sufficient condition to create “exceptions” to ATR harmony: the opaque generalization (ATR harmony) must still be obeyed on a “hidden” level (which is the level of the [±ATR] feature in this case) if exceptions to the surface realization of ATR harmony are to be created.

In this section, I have shown that the opaque interaction between ATR harmony and vowel assimilation in Yoruba can be accounted for in terms of differential classification. The crucial notion behind this account is the assumption that ATR harmony is a generalization over featural values, while full vowel assimilation is a generalization over segmental values. As has been explained above, this assumption is quite plausible given the nature of both processes. Because segmental and featural values are related to one another through violable constraints, the opaque interaction between ATR harmony and vowel assimilation can be explained as a mismatch between assimilated vowels and their [±ATR] classification, which occurs just in case this is triggered by the constraint which induces vowel assimilation.

3. Case study II: Bedouin Arabic

3.1 The facts: vowel epenthesis makes vowel raising opaque

Bedouin Arabic (Al-Mozainy 1981, McCarthy 2007b) is a language rich in opaque interactions. One of these opaque interactions takes place between the processes of vowel raising and high vowel epenthesis. In this section, I will provide an account of this interaction, as it has been described by McCarthy (2007b), in the framework of differential classification.

Vowel raising is a process which converts underlying short /a/ into a high vowel under a complex set of circumstances. The quality of short high vowels is mostly, if not completely[13], determined by the surrounding consonants. In non word-final open syllables, underlying /a/ is standardly realized as [i] (53a). However, there are two sets of specific contexts in which /a/ in a non-final open syllable is realized as [a] instead of [i]. These contexts are when the target vowel comes after a guttural consonant (53b), and when the target vowel is followed by either a coronal sonorant or a guttural consonant plus the vowel [a] (53c). On top of this set of complications, there is another exception: when the target vowel is followed by a syllable with a high vowel, then raising happens regardless of which consonant precedes the target vowel (53d).

(53)
a. *Default case: /a/ → [i] in a non-final open syllable*
/katab/ → [kitab] "he wrote"
/dafaʕ/ → [difaʕ] "he pushed"

b. *But: /a/ → [a] when adjacent to guttural*
/ʕabad/ → [ʕabad] "he worshiped"
/naʕas/ → [naʕas] "he dozed"

c. *But: /a/ → [a] when followed by coronal sonorant*
/sarag/ → [sarag] "he stole"
/manaʕ/ → [manaʕ] "he prohibited"

d. *However, overriding exceptions, always /a/ → [i] when followed by [i]*
/ħalim/ → [ħilim] "he dreamt"
/karih/ → [kirih] "he disliked"

Next to this process of raising, Bedouin Arabic also has a process of high vowel epenthesis between the members of certain, disallowed, consonant clusters. The vowel epenthesized is a short high vowel, which varies in quality between [i ~ u ~ ɨ] (depending on the quality of the surrounding segments). Epenthesis can occur in word-medial and word-final consonant clusters. In this account, I will abstract away from the quality of the short high vowel, as well as the quality of the segments in prohibited consonant clusters.

(54)
a. /jabr(u)kin/ → [ja<u>bur</u>kin] "they (she-camels) kneel"
/jasr(u)gin/ → [ja<u>sur</u>gin] "they (f.) steal"
b. /gabl/ → [gabil] "before"
/gabr/ → [gabur] "grave"

13 John McCarthy (p.c.).

As can be seen in the examples above, the instances of /a/ preceding the inserted high vowel are in non-final open syllable, and yet are realized as [a]. This means that epenthesis counterfeeds vowel raising: while [a] before a high vowel is normally always realized as a high vowel, this does not happen when the high vowel is created by epenthesis. I will account for this opaque interaction in terms of differential classification: it will be assumed that epenthetic high vowels, while being phonetically identical to underlying high short vowels, lack the feature [-high]. Because of this featural deficiency, epenthetic high vowels fail to trigger raising.

Section 3.2 will lay out an analysis of vowel raising, while section 3.3 provides an analysis of epenthesis; finally, section 3.4 will show how the opaque interaction between raising and epenthesis is derived.

3.2 Vowel raising

In the opaque interaction outlined above, vowel raising is the process which is made opaque. The facts of vowel raising are repeated below:

(55)
a. *Default case: /a/ → [i] in a non-final open syllable*
/katab/ → [kitab] “he wrote”
/dafaʕ/ → [difaʕ] “he pushed”

b. *But: /a/ → [a] when adjacent to guttural*
/ʕabad/ → [ʕabad] “he worshiped”
/naʕas/ → [naʕas] “he dozed”

c. *But: /a/ → [a] when followed by coronal sonorant*
/sarag/ → [sarag] “he stole”
/manaʕ/ → [manaʕ] “he prohibited”

d. *However, overriding exceptions, always /a/ → [i] when followed by [i]*
/ħalim/ → [ħilim] “he dreamt”
/karih/ → [kirih] “he disliked”

McCarthy (2007b) analyzes raising as a reduction process: the constraint which triggers raising is one which prefers shorter and less sonorous vowels (i.e., [i] is preferred over [a]) in “weak” syllables. The syllables that count as “weak” are those that are open and non-final[14]. This explains why, even if the segmentally driven alternations shown in (55b-d) above are factored out, non-final open syllables still show a general preference for [i] instead of [a] (which are supposed to be “weak”).

If the factor triggering raising is indeed the fact that /a/ is in an open and non-final syllable, then the problem in accounting for the opaque interaction between vowel raising and epenthesis (for example, /gabl/ → [gabil], *[gibil]; see preceding section) is that the [a] in /gabl/ → [ga.bil] is in a (non-final) open syllable and yet not raised (compare examples (55a,d) above). This means that a structural constraint (“no [a] in non-final open syllables”) which dominates faithfulness may suddenly be violated when some other process has taken place, which in turn necessitates an analysis in terms of sequential ordering of processes.

14 The full definition of “weak” syllables is non-final, open, and unstressed. However, the processes in Bedouin Arabic which apply to “weak” syllables (syncope and vowel raising) also occur in stressed syllables. McCarthy's analysis accounts for this by assuming that processes triggered in “weak” syllables derivationally precede stress assignment, so that the distinction between stressed and unstressed syllables is irrelevant for these processes.

However, I propose that vowel raising is triggered by a series of constraints which penalize the feature [-high] when it is followed by some other value of [±high]. I assume that every occurrence of [a] must also be classified as [-high], so this has the effect that every non-final [a] is penalized by one of these constraints. However, the crucial difference with McCarthy's (2007b) account is that the context which follows [a] matters for each individual constraint. The three constraints that I propose as triggers of raising are the following:

[-high,V][+high,V] ([-high][+high]): Assign one violation mark for every [-high] vowel which, on the consonant/vowel tier, is directly followed by a [+high] vowel within the same word.

[-high,V][-high,V] ([-high][-high]): Assign one violation mark for every [-high] vowel which, on the consonant/vowel tier, is directly followed by another [-high] vowel within the same word.

[-high,V][∅,V] ([-high][∅]): Assign one violation mark for every [-high] vowel which, on the consonant/vowel tier, is directly followed by vowel not associated with any value of [±high] within the same word.

The definition of this family of constraints presupposes that every vowel segment is associated with some feature, let us call it [+vocalic], which designates vowels as a class. I assume that every vowel, underlying or epenthetic, is required to have this feature in Bedouin Arabic, and that this is enforced by an undominated constraint (which I will omit from the explicit analysis).

The definitions above also presuppose appropriate assignment of the feature [±high]. I assume that every instance of [a] must be associated with the feature [-high] on the surface, which can be formalized by assuming that the constraint [a] → [-high] is undominated:

[a] → [-high]:

Assign one violation mark for every instance of [a] which is not associated with the feature value [-high].

The assignment of [+high] to high vowels is assumed to be less uniform (specifically, I assume that epenthetic high vowels may not be specified for [±high]); this behavior will be formalized later on.

This preference for [-high] vowels to be final in the word could have been expressed by a single constraint of the Align family (see McCarthy & Prince 1993), but the empirical facts of Bedouin Arabic suggest otherwise. As can be seen in (55) above, raising of low vowels before low vowels is blocked in certain conditions (55b,c), while raising of low vowels before underlying high vowels does not undergo blocking in these conditions (55d). This suggests that there are separate constraints against low vowels before high vowels (*[-high][+high]), and low vowels before low vowels (*[-high][-high]).

Also, as has been discussed above, low vowels fail to rise before epenthetic [i] (as opposed to underlying [i]). If it assumed that epenthetic vowels lack a [±high] specification, then the presence of raising before [i] vowels with a [+high] specification, and the lack of raising before [i] vowels without a [±high] specification can be explained by different rankings for the constraints *[-high][+high] and *[-high][∅]. The exact ranking of the three constraints with respect to other constraints in the language will be specified in the course of the remainder of this section.

I will first focus on the relative ranking of *[-high][+high] and *[-high][-high]. There are two conditions under which vowel raising before a low vowel is potentially blocked: immediately adjacent to a guttural consonant, or before a coronal sonorant.

The first of these conditions can be formalized by the following constraint:

*{Gi} = *{[+guttural],[+high]}: Assign one violation mark for every timing slot associated with [+guttural] directly preceding or following a timing slot associated with [+high]. (This constraint blocks raising when the target vowel is directly preceded or followed by a guttural consonant)

This constraint presupposes that there is a feature [±guttural], the positive value of which is assigned to every occurrence of a guttural consonant. Of course, the presence of this feature requires the presence of a series of high-ranked constraints which demand that each of the guttural consonants in this language is assigned the feature [+guttural]. To save space, I will take the presence of these high-ranked constraints for granted, and not make it explicit in the analysis.

The condition that raising does not occur before a coronal sonorant and a low vowel (i.e., in the context [__Ra]) can be phonetically explained as vowel harmony across coronal segments. It has been pointed out that vowel harmony across coronal segments is phonetically more plausible than across other types of consonants, and harmony across coronal segments only has been typologically attested in various languages (see Gafos & Lombardi (1999) and references therein).

However, I argue that this cannot be the synchronic grammatical interpretation of this phenomenon. This is because a constraint which promotes harmony across coronals is satisfied both by [a ... a] and [i ... i]. The configuration [a ... a] is penalized by the constraint *[-high][-high]. Because *[-high][-high], as will be seen, must be high-ranked (higher than faithfulness to vowels), and coronal harmony does not prefer [a … a] over [i … i], there is no way for a coronal harmony constraint to rule out an underlying form like /sara/ from turning into [siri] instead of [sara]. McCarthy's (2007b) analysis can use a coronal harmony constraint to enforce /sara/ → [sara] because the framework of that analysis, OT with Candidate Chains, only allows one single change to occur at every step of derivation, so that /sara/ may not be changed to [siri] in one single step, and [sara] is the only available candidate. However, the current analysis is couched in standard OT, which does not have this restriction on the GEN component, so that a coronal harmony constraint leads to the problem just mentioned. This will be shown in tableau form further on, when the ranking of *[-high][-high] will be formally established.

Because of this empirical problem, I argue that the synchronic representation of this particular condition under which vowel raising is blocked is simply a constraint which penalizes high vowels before a coronal sonorant:

*iR = *[+high][+corson]: Assign one violation mark for every [+high] segment followed by a coronal sonorant ([+corson]) segment.

Of course, this constraint, as the preceding one, presupposes the existence of a feature consistently assigned to a group of consonants. The presence of the feature [+corson] on every coronal sonorant of the language (or, equivalently, two features [+coronal] and [+sonorant] on every coronal sonorant) is assumed to be regulated by a series of undominated constraint, which will not be spelled out for the sake of brevity.

Both of these constraints must be ranked above *[-high][-high], because the lack of vowel raising promoted by these two constraints is chosen over the option with raising as promoted by *[-high][-high]. The following tableau shows that *{Gi} must dominate *[-high][-high]; *{Gi} would be violated if there were raising in the word [ʕabad].

(56)

[+G] [-hi] [-hi] / ʕ a b a d/	*{Gi}	*[-high][-high]
[+G] [-hi] [-hi] ʕ a b a d		*
[+G] [+hi] [-hi] ʕ i b a d	*W	L

The tableau below shows that *iR must be above *[-high][-high]: raising in the word [sarag] would lead to a violation of *iR.

(57)

[+cs] [-hi] [-hi] / s a r a g/	*iR	*[-high][-high]
[+cs] [-hi] [-hi] s a r a g		*
[+cs] [+hi] [-hi] s i r a g	*W	L

The constraint which triggers raising before a high vowel, *[-high][+high], in turn must be higher than the constraints *{Gi} and *iR, which block raising in their respective contexts. Raising before a high vowel is never blocked by either of these constraints, as evidenced by forms like /ħalim/ → [ħilim]: the winning candidate violates both *{Gi} and *iR constraints for the sake of not having a low vowel before a high vowel.

(58)

[+G] [-hi] [+hi] / ħ a l i m/	*[-high][+high]	*{Gi}	*iR
[+G] [+hi] [+hi] ħ i l i m		*	*
[+G] [-hi] [+hi] ħ a l i m	*W	L	L

Underlying /i/ is never changed into [a] when adjacent to a guttural. Therefore, the constraint *{Gi} must also be dominated by an NoChange constraint which prevents underlying [+high] values from being changed – this is shown below on the basis of the example /samiʕt/ → [simiʕt] "I heard" and the hypothetical mapping /gadir/ → [gidir]:

(59)

[+G] [-hi] [+hi] /s a m i ʕ t/	NoChange([+high])	*{Gi}
[+G] [+hi] [+hi] s i m i ʕ t		*
[+G] [+hi] [-hi] s i m a ʕ t	*W	L

(60)

[+cs] [-hi] [+hi] /g a d i r/	NoChange([+high])	*iR
[+cs] [+hi] [+hi] g i d i r		*
[+cs] [+hi] [-hi] g i d a r	*W	L

Thus, we have established a hierarchy of constraints which trigger vowel raising, so that only raising before a low vowel has exceptions, while raising in an open syllable before an underlying high vowel always takes place:

(61)

*[-high][+high]   NoChange[+high]

*{Gi} , *iR

*[-high][-high]

Vowel raising is a process which involves change rather than insertion or deletion, which means that the triggering constraints *[-high][+high] and *[-high][-high] must dominate NoChange constraints on [-high] values and [a] segments:

NoChange([-high]):

Assign one violation mark for every input instance of the feature value [-high] whose output correspondent on the [±high] tier is not identical with it[15].

15 Since features are seen as individual entities, they are also assigned their individual input-output correspondence relations, and NoChange constraints only examine the identity relation between an input and an output feature value rather than an input and an output segment which should share the same feature value.

NoChange([a]):

Assign one violation mark for every input instance of the segment [a] whose output correspondent on the segment tier is not identical with it.

The tableau below shows why NoChange([-high]) and NoChange([a]) must be dominated by *[-high][-high]. /katab/ surfaces as [kitab] because it is more important to avoid two consecutive non-high vowels than it is to be faithful to underlying [-high] and segmental values. Because *[-high][+high] must dominate *[-high][-high], it also dominates these two faithfulness constraints by transitivity.

(62)

[-hi] [-hi] /k a t a b/	*[-high][-high]	NoChange([-high])	NoChange([a])
[+hi] [-hi] k i t a b		*	*
[-hi] [-hi] k a t a b	*W	L	L

At the same time, vowel raising is categorically blocked in closed syllables. This can be formalized by a high ranking for positional variants of the faithfulness constraints shown above, which protect underlying [-high] and underlying [a] in closed syllables specifically:

NoChange([-high], Close-syll-head):

Assign one violation mark for every input instance of the feature value [-high] whose output correspondent on the [±high] tier is associated with the head of a closed syllable, but not identical with the input value.

NoChange([a], Close-syll-head):

Assign one violation mark for every input instance of [a] whose output correspondent on the segment tier is the head of a closed syllable, but is not identical with the input value.

The input /katabtum/ “you (m. pl.) wrote” surfaces as [kitabtum] because it is more important to maintain underlying [-high] and [a] in closed syllables than to avoid non-final non-high vowels. At the same time, open syllables are not affected by this faithfulness-over-markedness ranking, so that raising does occur in them.

If the input /katabtum/ is hypothesized to have no height feature in the closed syllable (which must be a possible input under the hypothesis of Richness of the Base), then it can be shown that NoChange([a], Closed-syll-head) must be ranked above the constraints which trigger raising. In the tableau below, I only consider candidates which obey the undominated constraint [a] → [-high] (which is not itself shown in the tableau).

(63)

[-hi] [+hi] /k a t a b t u m/	NoChange([a], Closed-syll-head)	*[-high][+high]	*[-high][-high]
[+hi] [-hi] [+hi] k i .t a b . t u m		*	
[+hi] [+hi] [+hi] k i .t i b . t u m	*W	L	
[-hi] [-hi] [+hi] k a . t a b . t u m		*	*W

The candidate in which the closed syllable [tVb] has a raised vowel is ruled out because it violates the high-ranked positional NoChangeity constraint, while the candidate in which there is no raising in the open syllable [kV] is excluded because it violates the triggering constraints more often than the (winning) candidate with raising in the open syllable only.

If, on the other hand, the input /katabtum/ is hypothesized to have no vowel segment in the syllable /tVb/, then it can be shown that NoChange([-high], Closed-syll-head) must dominate the two triggering constraints. Once again, only candidates which satsify undominated [a] → [-high] are considered[16].

(64)

[-hi] [-hi] [+hi] /k a t b t u m/	NoChange([-high], Closed-syll-head)	*[-high][+high]	*[-high][-high]
[+hi] [-hi] [+hi] k i .t a b . t u m		*	
[+hi] [+hi] [+hi] k i .t i b . t u m	*W	L	
[-hi] [-hi] [+hi] k a . t a b . t u m		*	*W

Note that it might be possible to obviate the violation of *[-high][+high] observed in the winning candidate in these two tableaux by lowering the final high vowel [u]: /katabtum/ → *[kitabtam]. Under the ranking established up until now, this candidate should win, since this candidate violates a lower-ranked markedness constraint than the winner: the desired winner violates higher-ranked *[-high][+high], while the candidate in which the last vowel is lowered only violates lower-ranked *[-high][-high] (this can be seen in the tableau below). Notice that the unattested lowering in /katabtum/ → *[kitabtam] is not blocked by constraints against change in a closed syllable, as such constraints have only been formulated for [+high] vowels.

The solution to the problem of lowering is to rank NoChange([+high]) – which has already been established as having a high ranking – above *[-high][+high]. In this manner, vowel lowering (rather than vowel raising) is never a solution to avoid violations of any of the constraints which trigger vowel raising.

16 I am assuming here that constraints against floating vowel features are undominated in Bedouin Arabic, so that floating vowel features in the input will lead to insertion of a corresponding vowel segment when this is allowed prosodically.

(65)

[-hi] [-hi] [+hi] /k a t a b t u m/	NoChange([+high])	*[-high][+high]	*[-high][-high]
[+hi] [-hi] [+hi] k i .t a b . t u m		*	
[+hi] [-hi] [-hi] k i .t a b . t a m	*W	L	*W

This ranking prevents the second vowel in a [-high][+high] configuration from being changed to a low vowel. However, there is another potential problem: the configuration of two non-high vowels in a row ([-high][-high]) could theoretically be avoided either by raising the first vowel, or by raising the second vowel. Bedouin Arabic only raises the first vowel in a pair, which needs to be accounted for. The answer to this problem can be found in the already established ranking *[-high][+high] >> [-high][-high]. When the second vowel in an /a...a/ sequence is raised, then the result violates the higher-ranked constraint *[-high][+high], whereas raising the first vowel in an /a...a/ sequence only results in no such violation, so that the candidate with raising in the second syllable is excluded.

This is demonstrated below on the basis of the mapping /katabna/ → [kitabna], *[kitabni] “we wrote”. In the pair of syllables [tab, na], the first vowel is prohibited from undergoing raising because it is in a closed syllable, but even in this circumstance, the second vowel is not allowed to raise to avoid the configuration [a...a].

(66)

[-hi] [-hi] [-hi] /k a t a b n a/	NoChange([-high], Cl-syl-hd)	*[-high][+high]	*[-high][-high]	NoChange([-high])
[+hi] [-hi] [-hi] k i t a b n a			*	*
[+hi] [+hi] [-hi] k i t i b n a	*!			**
[+hi] [-hi] [+hi] k i t a b n i		*!		**
[-hi] [-hi] [-hi] k a t a b n a			**!	

As can be seen in the tableau, the winning candidate violates *[-high][-high] once, because of the syllable pair [tab][na]. The candidate in which the closed syllable [tVb] undergoes raising is blocked by positional faithfulness. The candidate in which the [-high][-high] configuration is avoided by raising the final vowel of the word is excluded because it violates the higher-ranked constraint *[-high][+high]. Finally, the fully faithful candidate, which has no raising even in the first syllable of the word (which is open and followed by another value of [±high], and thus should always undergo raising) is excluded because it has more than one violation of *[-high][-high].

These considerations leave us with the following ranking of constraints (not counting undominated [a] → [-high]), which accounts for the segmental conditions on raising, as well as the condition that raising only occurs in non-final open syllables.

(67)

NoChange([-high], Cl-syll-hd) NoChange([a], Cl-syll-hd) NoChange([+high])

*[-high][+high]

*{Gi} CorHarmony

*[-high][-high]

NoChange([-high]) NoChange([a])

Having established this ranking of constraints to account for vowel raising, I will now illustrate how the properties of vowel raising are actually accounted for, through summary tableaux for the mappings /katabna/ → [kitabna], hypothetical /sara/ → [sara], /ʕabad/ → [ʕabad], and /ħalim/ → [ħilim][17].

(68)

[-hi] [-hi] [-hi] /k a t a b n a/	NoChange([-high/a], Cl-syll-hd)	NoChange([+high])	*[-high][+high]	*{Gi}	CorHarmony	*[-high][-high]	NoChange([-high/a])
[+hi] [-hi] [-hi] k i t a b n a						*	*
[-hi] [-hi] [-hi] k a t a b n a						**!	
[+hi] [+hi] [-hi] k i t i b n a	*!					*	**
[+hi] [-hi] [+hi] k i t a b n i			*!				**

In this tableau, it is shown that raising in open syllables followed by a low vowel must happen. In the second candidate, this is not the case, and this candidate is excluded because it violates *[-high][-high] (too many times). Raising in a closed syllable, as in the third candidate, is excluded by high-ranked positional faithfulness. Finally, raising in a final syllable, which could avoid violations of *[-high][-high], is excluded because it leads to a violation of higher-ranked *[-high][+high].

17 In these tableaux, the constraint pairs NoChange([-high], Cl-syll-hd) and NoChange([a], Cl-syll-hd), and NoChange([-high]) and NoChange([a]) will be collapsed.

(69)

[-hi] [-hi] [+cs] /s a r a /	NoChange ([-high/a], Cl-syll-hd)	NoChange ([+high])	*[-high] [+high]	*{Gi}	*iR	*[-high] [-high]	NoChange ([-high/a])
[-hi] [-hi] [+cs] s a r a						*	
[+hi] [-hi] [+cs] s i r a					*!		**
[-hi] [+hi] [+cs] s a r i			*!				
[+hi] [+hi] [+cs] s i r i					*!		****

In this tableau, the second candidate, in which the first vowel in an /a R a/ configuration is raised, violates the constraint *iR, which outranks any constraint violated by the faithful candidate. The third candidate, in which the second vowel in the /a R a/ configuration is raised, is excluded because it violates high-ranked *[-high][+high] – low vowels before high vowels are punished more severely than before other low vowels. These two candidates would also be excluded if *iR were replaced by a (phonetically more natural) constraint which promotes vowel harmony across coronal sonorants: both of these candidates would violate this constraint, which means that they would be excluded if the harmony constraint were ranked above *[-high][-high].

However, the fourth candidate, in which both vowels in the /a R a/ configuration are raised, would not be excluded by a harmony constraint. Whereas this candidate violates faithfulness to low vowel features and segments, while the desired winner (without raising) does not, it is also true that this fourth candidate has no violation of *[-high][-high], which is ranked above these faithfulness constraints. This means that this fourth candidate would win if there is no higher-ranked constraint excluding it. Because this candidate fully satisfies vowel harmony (or a constraint of the shape *aRa), a vowel harmony constraint would incorrectly predict that /a R a/ would map to [i R i]. This is why there must be a constraint of the shape *iR (i.e., against the sequence of a high vowel and a coronal sonorant), which, as shown in the tableau, excludes the candidate with raising of both vowels by virtue of being ranked above *[-high][-high].

(70)

[+G] [-hi] [-hi] / ʕ a b a d/	NoChange ([-high/a], Cl-syll-hd)	NoChange ([+high])	*[-high] [+high]	*{Gi}	*iR	*[-high] [-high]	NoChange ([-high/a])
[+G] [-hi] [-hi] ʕ a b a d						*	
[+G] [+hi] [-hi] ʕ i b a d				*!			*
[+G] [-hi] [+hi] ʕ a b i d	*!		*				*

This tableau shows how gutturals block vowel raising. The winning candidate violates *[-high][-high], violations of which are typically avoided through vowel raising. However, when the first vowel of this word is raised, as in the second candidate, this results in violation of the higher-ranked constraint *{Gi}. When the second vowel of the word is raised, this also results in the violation of a higher-ranked constraint: *[-high][+high] (no low vowel before a high vowel).

(71)

[+G] [+cs] [-hi] [+hi] / ħ a l i m/	NoChange ([-high/a], Cl-syll-hd)	NoChange ([+high])	*[-high] [+high]	*{Gi}	*iR	*[-high] [-high]	NoChange ([-high/a])
[+G] [+cs] [+hi] [+hi] ħ i l i m				*	*		*
[+G] [+cs] [-hi] [+hi] ħ a l i m			*!				
[+G] [+cs] [-hi] [-hi] ħ a l a m		*!				*	

This word has both conditions which would block vowel raising before a low vowel: a guttural adjacent to the low vowel, and a coronal sonorant following the low vowel. However, because the low vowel is followed by a high rather than a low vowel, raising still takes place. The second, fully faithful candidate illustrates why lack of raising is not an option in such cases: the constraint *[-high][+high] which this candidate violates is ranked higher than *{Gi} and *iR, which are violated when raising does take place in a guttural-and-coronal-sonorant environment. The third candidate shows that the configuration [a...i] may not be avoided by lowering the second vowel rather than raising the first vowel: vowel lowering is excluded by the high-ranked constraint NoChange([+high]). Therefore, the violations of *{Gi} and *iR which are assigned to the winning candidate – in which the first vowel is raised – are tolerated.

Now that the mechanism behind vowel raising has been explained, I will continue with an account of the process of epenthesis, after which I will show how vowel raising and epenthesis interact opaquely, and return to the function of the constraint *[-high][∅] which was introduced at the beginning of this section.

3.3 Epenthesis

Bedouin Arabic has a process by which a high vowel is inserted between certain pairs of consonants when they become adjacent. As has been shown at the beginning of this section, epenthesis enters into an opaque interaction with vowel raising: vowel raising only occurs before underlying high vowels, not before epenthetic high vowels. The essence of my account of this opaque interaction will be that epenthetic vowels in Bedouin Arabic lack a height feature, while vowel raising is conditioned by specifications on the vowel height tier.

The idea that epenthetic vowels may lack certain features and therefore not participate in certain processes is certainly not new. For instance, Piggott (1995) argues that epenthetic vowels, but not underlying, vowels in Mohawk lack morae, so that they are exempt from certain prosodic processes. Also, Prince & Smolensky (1993/2004) argue that epenthetic vowels are empty segmental positions which lack all featural content. However, such accounts, in which epenthetic and non-epenthetic vowels differ in their surface phonological specifications, have no explanation for why the epenthetic and non-epenthetic versions of a vowel have the same phonetic value[18].

Interestingly, it has been found by Gouskova & Hall (2009) that epenthetic and non-epenthetic [i] in Lebanese Arabic are indeed phonetically distinct for some speakers. However, other speakers of the same dialect have no acoustic distinction between epenthetic and non-epenthetic [i]. There is also no indication of such a distinction in Bedouin Arabic. In fact, the classification of the raising-epenthesis interaction in Bedouin Arabic as opaque depends on the assumption that epenthetic and non-epenthetic [i] are the same sound. If epenthetic and non-epenthetic [i] were not the same sound, then the generalization that low vowels always raise before “true” (non-epenthetic) [i] would be surface-true and surface-apparent: therefore, not opaque.

The account that will be offered in this section will account for both pieces of the puzzle: epenthetic and non-epenthetic [i] function differently at the featural level (only non-epenthetic [i] triggers raising), but they have the same segmental value (meaning that they will be pronounced the same). This is possible because the featural and segmental level are independent levels of description of the same utterance. Since it is the segmental tier that lower-level (phonetic) generalizations pay attention to, segmental equivalence of underlying and epenthetic high vowels will lead to their phonetic equivalence.

Epenthesis in Bedouin Arabic occurs whenever certain kinds of two-consonant sequences threaten to become tautosyllabic. The exact set of disallowed sequences is not important here: only the prosodic factors are relevant. The vowel epenthesized between the two consonants is a high vowel ([i~u][19]):

18 A possible solution to such problems it to relegate them to phonetic implementation; however, there does not seem to be a general theory of phonetic implementation which specifies to what extent phonetic implementation rules may lead to categorical neutralization of a phonological contrast.

19 High short vowels in Bedouin Arabic fluctuate between the values [i ~ u ~ ɨ], depending on their segmental environment. It is not clear whether the language has a definite phonemic contrast betwen short /i/ and /u/ (John McCarthy, p.c.). To simplify the discussion of epenthesis, I will assume that there is no contrast between (etymologically distinct) /i/ and /u/.

(72)
a. /jabr(u)kin/ → [jaburkin] "they (she-camels) kneel"
/jasr(u)gin/ → [jasurgin] "they (f.) steal"
b. /gabl/ → [gabil] "before"
/gabr/ → [gabur] "grave"

The examples in (72b) above show that an [a...i] configuration created by epenthetic [i] does not lead to raising of the first vowel ([a]) in that configuration – which means that epenthesis fails to feed, and therefore counterfeeds, raising. It is this opaque interaction that I will account for by differential classification in this section.

Let us assume that epenthesis is triggered by a constraint Cluster-Condition, which is a shorthand for :

Cluster-Condition:
Assign one violation mark for every two-consonant cluster (of the specified type), whenever both of its members are part of the same syllable.

As has been said before, I propose that the vowel epenthesized between the two consonants of an offending cluster has a segmental value (and perhaps some feature values like [+vocalic]), but no value for [±high]. Thus, the winning candidate for a word like /gabl/ → [gabil] would have a representation like the following:

(73)
[-hi]
|
g a b i l

To allow this "featural underspecification" of epenthetic [i] specifically, the constraints Dep([i]) – which militates against insertion of the segment [i] – and [i] → [+high] – which militates against [i] without a [+high] feature – must be ranked below Cluster-Condition, because the insertion of [i] without a [+high] feature is triggered by Cluster-Condition.

Dep([i]): Assign one violation mark for every output instance of [i] which has no correspondent in the input.

(74)

[-hi] /g a b l/	Cluster-Condition	Dep([i])	[i] → [+high]
[-hi] g a b i l		*	*
[-hi] g a b l	*W	L	L

Furthermore, the feature [+high] must be prevented from being inserted on an epenthetic [i], which would lead to raising of the preceding vowel, which means that the constraint against inserting [+high] – Dep([+high]) – should dominate [i] → [+high]:

(75)

[-hi] /g a b l/	Dep([+high])	[i] → [+high]
[-hi] g a b i l		*
[+hi] [+hi] g i b i l	*W	L

Finally, the constraint against inserting the segmental value [a] – Dep([a]) – must be ranked above Dep([i]) to ensure that it is better to insert [i] than it is to insert [a] in an epenthetic context[20]. Also, Dep([a]) must dominate the constraint against [i] without a height feature: [i] → [+high]. This is because epenthesis of [a] instead of [i] avoids a violation of this constraint.

(76)

[-hi] /g a b l/	Dep([a])	Dep([i])	[i] → [+high]
[-hi] g a b i l		*	*
[-hi] [-hi] g a b a l	*W	L	L

Thus, we obtain the following summary ranking of constraints necessary for epenthesis:

(77)

Cluster-Cond, Dep([a])　　　Dep([+high])

Dep([i])　　　[i] → [+high]

The summary tableau below shows how this ranking derives obligatory epenthesis of [i] which is not associated with a [+high] feature. The result that epenthetic vowels must have no value for [±high] will be crucial in my account for the opaque interaction between vowel raising and epenthesis.

20 Strictly speaking, epenthesis of [a] could be prevented by either Dep([-high]) or Dep([a]). However, I will argue later on that Dep([-high]) must be low-ranked.

(78)

[-hi] /g a b l/	Cluster-Cond	Dep([+high])	Dep([a])	[i] → [+high]	Dep([i])
[-hi] g a b i l				*	*
[-hi] g a b l	*!				
[+hi] [+hi] g i b i l		*!			*
[-hi] [-hi] g a b a l			*!	*	

The candidate without epenthesis violates high-ranked Cluster-Cond. The candidate in which a [+high] feature is inserted is penalized by high-ranked Dep([+high]). Finally, the candidate in which [a] instead of [i] is inserted is ruled out by Dep([a]). Thus, the winning candidate is justified in violating the low-ranked constraints [i] → [+high] and Dep([i]).

The hierarchy just established to account for vowel epenthesis is completely disjunct with the hierarchy established to account for vowel raising (see (67) in the previous subsection). To connect the two subhierarchies, it is essential to consider inputs that contain vowel segments without a height feature.

One of the most important questions given the principle of Richness of the Base (Smolensky 1996) is why epenthetic vowels, which are not associated with a height feature, create surface exceptions to vowel raising (i.e., allow the existence of [a...i] sequences), but underlying [i] with no height feature may not lead to exceptions to vowel raising. The latter can be inferred from the fact that vowel raising has no exceptions other those created by epenthesis, and the fact that Richness of the Base demands that every possible combination of representational elements, when used as an input, leads to a grammatical output.

I assume that inputs of the form /C [a, -high] C [i] C/, with a high vowel without a height feature following a low vowel, do not come out unchanged, but, instead, end up with raising of the low vowel, and lowering of the high vowel.

Thus, an input such as

  [-hi]
/t a l i k/

will surface as

  [+hi] [-hi]
[t i l a k].

Notice that the "lowering" in this case does not involve a change from [+high] to [-high], so that it is not excluded by high-ranked Dep([+high]).

The mapping shown above can be optimal in the current system because each representational element is seen as autonomous, and thus, both features and segments have both NoChange and Maximality/Dependency constraints protecting them. In this case, the mapping shown above can be obtained by ranking the NoChange constraints on underlying /a/ and /i/ below Dep([+high]) and [i] → [+high].

NoChange([a]):

> Assign one violation mark for every input instance of [a] whose counterpart on the segmental tier is not identical to it.

NoChange([i]):

> Assign one violation mark for every input instance of [i] whose counterpart on the segmental tier is not identical to it.

Dep([+high]) and [i] → [+high] must also dominate two faithfulness constraints referencing the feature value [-high]: NoChange([-high]) and Dep([-high]). This is because I have assumed that every instance of [a] in Bedouin Arabic must have the feature [-high] associated with it, so that underlying [-high] over the first vowel in the current input candidate must change to [+high], while [-high] must be inserted to allow the second vowel to change to [a].

(79)

[-hi] /t a l i k/	Dep([+high])	[i] → [+high]	NoChange([a])	NoChange([i])	NoChange([-high])	Dep([-high])
[+hi] [-hi] t i l a k			*	*	*	*
[-hi] t a l i k		*W	L	L	L	L
[+hi] [+hi] t i l i k	*W		*	*	*	L

The fully faithful candidate (with height-feature-less [i]) is ruled out by [i] → [+high]. The candidate in which the second syllable of the word has [i] with a [+high] feature is excluded by higher-ranked Dep([+high]). The remaining possibility is to change [i] into [a], which is the winning possibility[21]. The high-ranked positional faithfulness constraints protecting vowels in closed syllables – NoChange([a], ClosedSyllable) and NoChange([-high], ClosedSyllable) – are not shown in this tableau, but these constraints could not prevent [i] from changing into [a], even if it is in a closed syllable, because these constraints only protect underlying low vowels.

Since I have assumed that the constraint [a] → [-high] is always obeyed in the language, it must also dominate the conflicting constraint Dep([-high]): it is justified to insert the feature [-high] onto underlying [a] without an underlying height feature. If it is assumed that underlying [a, ∅high] ([a] without a height feature) normally maps onto [a, -high], then Dep([-high]) must also be dominated by NoChange([a]), because satisfaction of [a] → [-high] is possible either through insertion of [-high] or through change of the segmental value. This is illustrated below:

21 I have not included candidates which violate the raising generalizations ([-high] is not allowed in open syllables before another [±high] value); such candidates would avoid violations of NoChange([a]) and NoChange([+high]), but since it has already been established that these constraints must be dominated by raising-triggering constraints, it is clear that such candidates would not win.

(80)

/b a/	[a] → [-high]	NoChange([a])	Dep([-high])
[-hi] b a			*
b a	*!		
b i		*!	

Another interesting class of potential inputs is one in which the last vowel is /a/ without a height feature. Under the current ranking, such inputs would lead to potential exceptions to the raising generalizations. This is shown in the tableau below for hypothetical /t [a,∅high] b [a,∅high] k/:

(81)

/t a b a k/	[a] → [-high]	Dep([+high])	[i] → [+high]	Dep([-high])
[-hi] [-hi] t a b a k				**
[-hi] t i b a k			*!	*
[-hi] t a b i k			*!	*
[-hi] t a b a k	*!		*	*
[+hi] [-hi] t i b a k		*!		*
[+hi] [+hi] t i b i k		*!*		

The undesirable winning candidate in this tableau is one which is an exception to raising: /tabak/ → [tabak]. Candidates in which one of the vowels is raised to [i] without a height feature are excluded by [i] → [+high]. Candidates in which surface [a] is not always associated with [-high] are excluded by undominated [a] → [-high]. Finally, candidates in which vowels are raised to [i] with a [+high] feature are excluded by high-ranked Dep([+high]).

The problematic candidate [t [a,-high] b [a,-high] k] which won in the tableau above can be excluded by ranking the raising-triggering constraint *[-high][-high] above [i] → [+high]:

(82)

/t a b a k/	*[-high][-high]	[i] → [+high]
[-hi] t i b a k		*
[-hi] t a b i k		*
[-hi] [-hi] t a b a k	*W	L

However, of the two remaining candidates, [tibak] and [tabik], only one is desirable: [tibak], since it obeys the raising generalizations. This is where the third raising-triggering constraint, *[-high][∅], which was mentioned at the beginning of the previous subsection, becomes important. This constraint prefers configurations in which [-high] is maximally final, even if it is not followed by other values of [±high]. Because of this, the constraint will prefer [tibak], where [-high] is absolutely final in the word, to [tabik]:

(83)

/t a b a k/	*[-high][∅]	[i] → [+high]
[-hi] t i b a k		*
[-hi] t a b i k	*W	*

Note that the constraint *[-high][∅] is violated by cases of epenthesis like /gabl/ → [gabil]. This means that some faithfulness constraint(s) must prevent such mappings from undergoing raising (or some other strategy to prevent [-high] from preceding a height-feature-less vowel).

I assume that the faithfulness constraints that prevent this from happening are NoChange([a]) and NoChange([+high]), the same ones violated when raising takes place. If we assume, for the moment, that the underlying form of [gabil] only has a segmental value [a] and no underlying [-high] feature, then raising in this case can be prevented by ranking NoChange([a]) above *[-high][∅]:

(84)

/g a b l/	NoChange([a])	*[-high][∅]
[-hi] g a b i l		*
g i b i l	*W	L

If, on the other hand, it is assumed that the underlying form [gabil] only has a feature value [-high] and no underlying segmental value [a], then the parallel argument can be made for NoChange([-high]) >> *[-high][∅][22].

22 I assume that Bedouin Arabic does not allow floating [±high] features.

(85)

[-hi] /g　　b l/	NoChange([-high])	*[-high][∅]
[-hi] g a b i l		*
[+hi] g i b i l	*W	L

The ranking NoChange([-high]), NoChange([a]) >> *[-high][∅] still derives the correct result for the input /tabak/ without height features, which is shown in the tableau below:

(86)

/t a b a k/	*[-high][-high]	[i] → [+high]	NoChange([a])	*[-high][∅]
[-hi] t i b a k		*	*	
[-hi] t a b i k		*	*	*!
[-hi] [-hi] t a b a k	*!			

Because the winning candidate and the counter-candidate [tabik] have the same violation pattern for [i] → [+high] and NoChange([a]), the lower-ranked constraint *[-high][∅] is able to distinguish between the two candidates and designate [tibak] as the winner. The candidate [tabak], which violates neither [i] → [+high] nor NoChange([a]), is excluded by its violation of higher-ranked *[-high][-high].

The ranking that has now been arrived at is summarized in the following diagram:

(87)

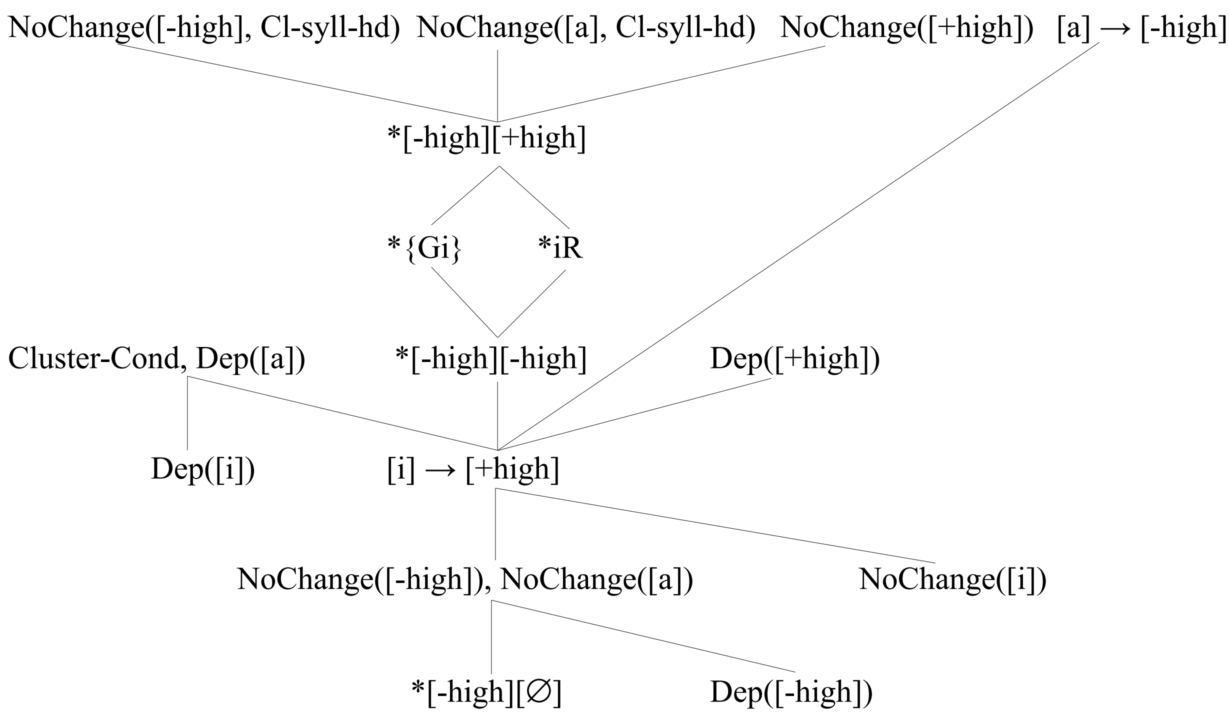

3.4 Explaining the opaque interaction

As has been briefly outlined at the beginning of this subsection, the opaque interaction between raising and epenthesis, as in /gabl/ → [gabil], *[gibil], arises through the required absence of a [+high] feature on the epenthetic vowel. Although all examples of epenthesis given here are opaque, the question of opacity has not been addressed explicitly. The constraint ranking as given above is, in fact, sufficient to account for the opaque interaction of epenthesis and raising – and the tableau below shows how (only crucial constraints are shown):

(88)

[-hi] /g a b l/	*[-high][+high]	Cluster-Cond	Dep([+high])	Dep([a])	[i] → [+high]	Dep([i])	NoChange ([-high])([a])	*[-high][∅]
[-hi] g a b i l					*	*		*
[-hi] g a b l		*!						
[+hi] [-hi] g i b a l				*!			*	
[-hi] [+hi] g a b i l	*!		*			*		
[+hi] [+hi] g i b i l			*!			*	*	
[+hi] g i b i l					*	*	*!	

As demonstrated by the second candidate, lack of epenthesis is ruled out by high-ranked Cluster-Cond. Epenthesis of [a] instead of [i] is excluded by its violation of Dep([a]).

Giving the epenthetic vowel the feature [+high] violates *[-high][+high], the constraint which triggers raising. However, if the epenthetic vowel has the feature [+high] and there is also raising of the preceding vowel (as in the fourth candidate, in which there is a transparent interaction of raising and epenthesis), then there is still a violation of high-ranked Dep([+high]).
Finally, if there is raising of the non-epenthetic vowel but no insertion of [+high] on the epenthetic vowel (to avoid a violation of Dep([+high]), then this candidate is harmonically bounded by the winning candidate without raising – it has a faithfulness violation for no reason, because the lack of a [+high] feature on the epenthetic [i] makes the constraint *[-high][+high] non-applicable to this candidate. Thus, the opaque candidate [gabil] with a [-high] feature on [a] remains as optimal.

The same ranking of constraints does produce vowel raising when [i] is not epenthetic; this is shown below on the basis of /karih/ → [kirih]. Once again, only crucial constraints are shown.

(89)

[-hi] [+hi] /k a r i h/	NoChange ([+high])	*[-high] [+high]	Cluster-Cond	Dep ([+high])	[i] → [+high]	*[-high] [-high]	Dep([i])	NoChange ([-high])/ ([a])
[+hi] [+hi] k i r i h								*
[-hi] [+hi] k a r i h		*!				*		
[-hi] k a r i h					*!			
[-hi] [-hi] k a r a h	*!							

Lack of raising in the name of faithfulness, as in the second candidate, is punished by high-ranking *aCi. Lack of raising by virtue of not having a height feature on the vowel [i] is not allowed, because it violates [i] → [+high]. Finally, lowering of the second [i] instead of raising of the first [a] is ruled out because this result, while avoiding a violation of *aCi, runs into a violation of *NoChange([+high]), which is ranked higher than NoChange([-high]). In this manner, raising is the optimal outcome.

In this section, I have explained how the opaque interaction between vowel raising and vowel epenthesis is accounted for in a differential classification framework. Epenthetic high vowels have the same segmental value as non-epenthetic high vowels, yet epenthetic high vowels are prevented from having a height specification, while non-epenthetic high vowels are forced to have a height specification. Because it is the height specification of a high vowel which triggers the raising of a preceding low vowel, epenthetic high vowels are correctly predicted to not trigger the raising of low vowels preceding them.

4. Discussion and comparison with other theories

The preceding sections of the paper have presented the idea of differential classification, and how this idea works to derive cases of phonological opacity in standard OT without adding any extensions to its basic architecture. However, presenting a new framework of analysis does not necessarily mean that a meaningful result has been reached: this framework may simply be a notational variant of an earlier analytical system, or it may not add any new predictions or connections to our understanding of a phenomenon while only adding more complexity to its conceptual structure.

In this section, I will argue that differential classification introduces both conceptual innovation and new empirical predictions. The conceptual innovation consists in a new, reductionist hypothesis regarding the synchronic representation of opaque interactions: opacity is argued to arise from the existence of different levels of taxonomy in the representation of sound systems, a concept which can also explain other phonetic and phonological phenomena. Differential classification makes predictions for both the typology of opaque interactions across languages, and for experimental work on opacity. These predictions differ from the predictions made by other theories of opacity.

In will present these conceptual and empirical projections in the form of comparison with three other theories: Stratal OT (Bermúdez-Otero 1999, Kiparsky 2000), OT with Candidate Chains (McCarthy 2007), and extensions of Containment Theory (Prince & Smolensky 1993/2004, Oostendorp 2004, 2007, 2008). Section 4.1 will compare differential classifications with these three approaches on a conceptual level, while section 4.2 will compare the empirical predictions made by these approaches.

4.1 Comparison with other theories: conceptual innovation

According to differential classification, phonological opacity may arise whenever two generalizations appeal to representations of the same sound(s) on different taxonomic levels, and the constraints relating these representations to one another are low-ranked. Sections 2 and 3 have shown in detail how such scenarios account for opaque interactions.

In this manner, the only basic concept required to account for both opaque and transparent interactions between processes is the presence of multiple representations of the same sound on different taxonomic levels which are related in a non-trivial way. I will show below that the concept of various taxonomic levels in the representation of sound structure has also been invoked by other work that deals with various, disparate phenomena not directly related to opacity. In this way, it can be argued that differential classification works towards a unified explanation for various phenomena.

However, before making argument, I will compare this conceptual motivation for opacity with the concepts that underlie the other theories mentioned above: Stratal OT, OT with Candidate Chains (OT-CC) and Containment Theory-based approaches.

4.1.1 Stratal OT

According to Stratal OT, phonological optimization interacts cyclically with morphosyntax: bare morphological stems pass through one layer of (OT) phonological grammar (the stem-level grammar); the outputs of this phonological grammar undergo morphological processes to form complete words, and then pass through another layer of phonological grammar (the word-level grammar). Finally, the outputs of word-level phonology undergo syntactic processes to combine into phrases, and these phrases pass through a final component of phonological grammar: phrase-level phonology. This is illustrated by the schema below:

(90)

stems

↓

stem-level phonology

↓

morphology

↓

word-level phonology

↓

syntactic rules

↓

phrase-level phonology

↓

output

The various "strata" of phonological grammar in this model - stem-level, word-level and phrase-level phonology - are assumed to be grammars with the same constraints but (potentially) different rankings of these constraints. From this setup, it follows that one and the same sound form can be subjected to different constraint rankings in different orders.

For instance, given the Canadian Raising example discussed in the introduction (section 1.3), if the word-level grammar has the ranking Faith-voice >> *VtV (no intervocalic [t]) and the phrase-level grammar has the ranking *VtV >> Faith-voice, then the output of word-level phonology will contain intervocalic [t], but not the output of phrase-level phonology[23]:

(91) Word-level phonology: /bɛtɪŋ/ → [bɛtɪŋ]

/bɛtɪŋ/	Faith-voice	*VtV
bɛtɪŋ		*
bɛd̺ɪŋ	*	

(92) Phrase-level phonology: [bɛtɪŋ] → [bɛd̺ɪŋ]

/bɛtɪŋ/	*VtV	Faith-voice
bɛtɪŋ	*	
bɛd̺ɪŋ		*

In the first tableau, which stands for the word-level stratum of phonology, the input /bɛtɪŋ/ is faithfully preserved in the output, because Faith-voice is ranked higher than *VtV. However, the phrase-level grammar, as shows in the second tableau, has the opposite ranking of these two constraints, causing the flapped variant [bɛd̺ɪŋ] to win in the phrase-level stratum of phonological grammar.

Because there is an intermediate, non-pronounced, output of at which intervocalic underlying /t/s are actually voiceless, there is a level at which the opaque generalization of Canadian raising is true (i.e., raising occurs only and always before voiceless obstruents): this level is the output of stem-level phonology. Thus, if the constraint ranking which triggers Canadian raising (CanRaising >> Faith-low) occurs in stem-level phonology, while word-level phonology requires faithfulness to diphthong quality (Faith-low >> CanRaising), the opaque interaction can be obtained. This is shown in the tableaux below:

CanRaising: Assign one violation mark for every instance of [aɪ] before a voiceless consonant, or [əɪ] before a voiced consonant.

23 The current analysis deviates from, and simplified with respect to the analysis in Bermúdez-Otero (2003, 2004).

(93) Word-level phonology: /ɹaɪtəɹ/ → [ɹəɪtəɹ]

/ɹaɪtəɹ/	Faith-voice	CanRaising	*VtV	Faith-low
[ɹaɪtəɹ]		*!	*	
[ɹəɪtəɹ]			*	*
[ɹaɪd̥əɹ]	*!			
[ɹəɪd̥əɹ]	*!	*		*

(94) Phrase-level phonology: [ɹəɪtəɹ] → [ɹəɪd̥əɹ]

/ɹaɪtəɹ/	*VtV	Faith-low	Faith-voice	CanRaising
[ɹaɪtəɹ]	*!	*		*
[ɹəɪtəɹ]	*!			
[ɹaɪd̥əɹ]		*!	*	
[ɹəɪd̥əɹ]			*	*

In the first of these two tableaux, the one corresponding to word-level phonology, the constraints Faith-voice and CanRaising are ranked high. This means that the first, fully faithful candidate is ruled out by high-ranked CanRaising - for it has an unraised diphthong before a voiceless consonant. At the same time, the two last candidates are ruled out by high-ranked Faith-voice, because these candidates have changed the voicing value of the medial consonant in this word. This means the the second candidate, which has raising but no flapping, wins.

In the second tableau, however - which corresponds to phrase-level phonology - the constraints *VtV and Faith-low are high-ranked. This means that the first two candidates, which do not have flapping, are ruled out because they violate the constraint *VtV. The third candidate, which has flapping and a non-raised diphthong, is ruled out by high-ranked Faith-low: the raised diphthong which was mandated in word-level phonology now becomes the standard to which phrase-level phonology must be faithful. Thus, the last candidate of the second tableau, a candidate with both flapping and diphthong raising, becomes the winner of the last stratum of phonological grammar.

Other opaque interactions can be obtained in a similar fashion: opaque generalizations are true of an “earlier” stratum of phonology than the generalizations that make them opaque (see, for instance, Bermúdez-Otero 1999:chapter 3.3.2 for another example of an analysis of an opaque phenomenon in Stratal OT). Of course, phonological generalizations are not randomly assigned to the various strata of phonological grammar: instead, the assignment of a phonological generalization to a particular stratum is synonymous to that phonological generalization applying to a certain type of morphosyntactic domain. Thus, if a generalization is enforced by the stem-level grammar but not other strata, then this generalization applies to stems only. Similarly with generalizations that are enforced up to the word-level grammar: these generalizations apply to words, but not phrases.

Given these considerations, one may say that the reason why an interaction between phonological generalizations can be opaque in the Stratal OT model is that the opaque generalization applies to a smaller morphosyntactic domain than the generalization that makes opaque. This is somewhat similar to the explanation that differential classification gives to opaque interactions: the opaque generalization must refer to a higher taxonomic level of phonological representation than the generalization that makes opaque. The similarity between these two hypotheses is the claim that the possibility for two generalizations to interact opaquely depends on the inherent nature of the generalizations themselves.

However, the aspects of generalizations that enable opaque interaction in Stratal OT versus in the differential classification view are different. The size of morphosyntactic domains is clearly related to taxonomic levels (stems may be contained in various words, words can be combined into many different phrases), but the differential classification view claims that it is not the taxonomic levels of morphosyntax, but the taxonomic levels of phonology itself that matter for opacity. This important conceptual difference results in different empirical predictions with regard to the range of possible opaque interactions, as will be seen in section 4.2.3.

4.1.2 OT with Candidate Chains

OT with Candidate Chains (OT-CC; McCarthy 2007b) is a theory in which serial ordering of phonological processes is re-introduced, after it had been eliminated from standard OT (see Prince & Smolensky 1993/2004). More specifically, the GEN component of OT is hypothesized to create output candidates by changing inputs one step at a time, instead of simply mapping an input string to every possible output string, however divergent from the input. These micro-changes performed by GEN, known as Localized Unfaithful Mappings, are only allowed to take place if they improve the overall harmony of the form given the constraint ranking of the language - in other words, GEN may only make a certain change at a particular moment in the derivation when this change somehow makes the form better for the grammar.

The EVAL component of the grammar is also slightly different from EVAL in standard OT: it has access to the exact sequence of Localized Unfaithful Mappings (LUMs) that was needed to create each output candidate. Each LUM (Localized Unfaithful Mapping) involves one violation of a Faithfulness constraint at a particular place in the word. Therefore, sequences of LUMs are expressed as sequences of Faithfulness constraint x location combinations (for instance: <Ident(low)@2, Ident(voice)@3>, which translates to “first, Ident(low) was violated at the second segment of the word, and then, Ident(voice) was violated at the third segment of the word”).

Because the fact that output candidates have such sequences of LUMs attached to them, EVAL has access to the serial derivation from input to output candidate. Because of this, there can be constraints on how this serial derivation may proceed. Such constraints are called Precedence (Prec) constraints: they specify that certain kinds of unfaithful mappings must always precede other kinds unfaithful mappings.

The concept of Prec(edence) constraints is similar to extrinsic rule ordering in rule-based phonology, according to which rule A may be restricted to applying only before rule B. Just as the concept of rule ordering can account for opaque interactions, restrictions on which unfaithful mappings may follow one another can give rise to opacity. If a high-ranked Prec-constraint says that all unfaithful mappings of type X must precede all unfaithful mappings of type Y, it can block a certain process A that involves a mapping of type X from repairing a configuration created by process B, if B involves a mapping of type Y.

For instance, if Canadian Raising involves changing the [±low] value of a diphthong nucleus, which violates the faithfulness constraint Ident(low), while the process of flapping requires changing the voicing value of a stop, which violates Ident(voice), then the opaque interaction between raising and flapping can be expressed through a high ranking of the constraint Prec(Ident(low), Ident(voice)). In informal terms, this constraint demands that every violation of Ident(low) is incurred before every violation of Ident(voice).

The tableau below shows Prec(Ident(low),Ident(voice) as dominating the Markedness constraint CanRaising (which demands that raised diphthongs precede voiceless consonants, and unraised diphthongs precede voiced consonants). CanRaising is also dominated by the Markedness constraint *VtV (which prohibits voiceless [t] in the Canadian Raising contexts). These Markedness constraints, in turn, dominate Ident(low) and Ident(voice), respectively.

The input for this tableau has the raised diphthong [əɪ] rather than the non-raised diphthong [aɪ].

(95)

/ɹəɪtɚ/	Prec(Id(low), Id(vce))	*VtV	CanRaising	Ident(low)	Ident(voice)
ɹəɪd̬ɚ <Id(vce)@3>			*		*
ɹaɪd̬ɚ <Id(vce)@3, Id(low)@2>	*!			*	*
ɹəɪtɚ <Id(low)@2>		*!			
ɹaɪtɚ <Id(vce)@3>		*!	*	*	

The second candidate in this tableau has first undergone flapping (voicing change) and then lowering of the raised diphthong from [əɪ] to [aɪ]: /ɹəɪtɚ/ → [ɹəɪd̬ɚ] → [ɹaɪd̬ɚ]. The inverse order of operations would not be possible, because lowering of [əɪ] to [aɪ] at the stage where /t/ has not yet changed to [d̬] would introduce a violation of CanRaising (as unraised diphthongs before voiceless consonants violate CanRaising): /ɹəɪtɚ/ *→ [ɹaɪtɚ] (→ [ɹaɪd̬ɚ]). Because GEN may only perform changes to the input which remove, rather introduce, violations of high-ranked constraints, the order <lowering, flapping> cannot be generated and thus never enters the set of outputs to be evaluated by EVAL.

This second candidate (which has a transparent interaction of raising and flapping) were to win in this tableau if it were not for the Prec constraint: this candidate satisfies both high-ranked Markedness constraints. However, because the (high-ranked) Prec constraint is violated by this candidate (which is the only possible candidate with an unraised vowel and a flapped consonant given this input), there is no possibility that the second candidate would win. The third and fourth candidates, which have no flapping, are ruled out by their violation of *VtV, which dominates CanRaising. Because of this, the opaque mapping /ɹəɪtɚ/ → [ɹəɪd̬ɚ] wins.

Thus, Prec constraints can eliminate certain transparent interactions of processes through penalizing a particular sequence of their application. Two generalizations may only interact opaquely if there is a high-ranked constraint which prevents the unfaithful mapping involved in the opaque generalization from occurring after the unfaithful mapping involved in the generalization that makes opaque.

This is very different from the factor that differential classification hypothesizes as the one underlying opacity: opaque interactions arise from inherent properties of generalizations, not from restrictions on the order of operations needed for these generalizations. As a consequence, the empirical predictions made by these two contrasting views are quite different - some of these differences in empirical predictions will be spelled out in section 4.2.3 below.

4.1.3 Containment-based approaches

One other approach to phonological opacity in OT has a representational rather than derivational character, and is based on the idea that the input to each OT computation is fully present in every output candidate made: whenever input and output diverge, the input elements that are not pronounced in the output are still covertly present in the representation of the output. This idea (called Containment Theory) was originally formulated in Prince & Smolensky (1993/2004) as the first theory of Faithfulness in OT. In later work, Containment Theory was abandoned in favor of Correspondence Theory (McCarthy & Prince 1994), in particular because the problems the original formulation of Containment Theory with epenthesis. However, recent work (Oostendorp 2004, 2007, 2008, Revithiadou 2007) has resurrected Containment Theory in a revamped formalism which avoids previous pitfalls.

According to this more recent version of the framework, called Colored Containment, each phonological object consists of a string of root nodes cum features, some of which are parsed by morphological structure, and some of which are parsed by prosodic structure. The elements that are parsed by morphological structure together form the underlying representation, while the elements that are parsed by prosody together form the pronounced surface form. When an underlying element is deleted or changed, it will be parsed by morphological but not prosodic structure. When a surface element is epenthetic, it will be parsed by prosodic but not morphological structure.

(96) a. No change between input and output: /matao/ → [matao][24]

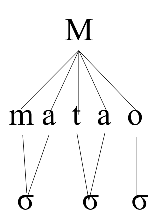

b. Deletion: /matao/ → [mato]; /a/ is not parsed by prosodic structure

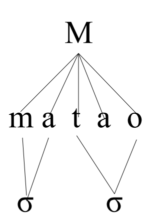

24 In the following diagrams, the symbol M stands for “morpheme”.

c. Epenthesis: /matao/ → [matawo]; [w] is not parsed by morphological structure

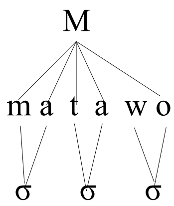

This model of phonological representation gives the phonological grammar access to the divergence between input and output (in a fashion somewhat similar to OT-CC). Every element in a Markedness constraint may be specified as having a certain affiliation (morphological or prosodic); this means that some constraints may not see every instance of a certain segment: for instance, if a constraint only sees prosodically parsed [t], it will ignore all other instances of [t] in the representation.

Because different constraints see different parts of the representation, different constraints may have contradictory views of one and the same part of the representation. Take, for instance, the case of Canadian Raising. The constraint against [t] in relevant contexts could be formulated such that only prosodically parsed [t] is prohibited in between two prosodically parsed vowels: *[VtV]$_{pros}$. On the other hand, the constraint against prosodically parsed diphthongs [aɪ, aʊ] before voiceless segments may be formulated without the requirement that the voiceless segment in question be parsed by prosodic structure: *[aɪ, aʊ]$_{pros}$[-voice]. In a word like “writer”, underlying /t/ will be present in every output candidate made from that input. Thus, the forms in (97a,b) are possible output candidates of “writer”, but (97c) is not, because underlying /t/ is not present.

(97)

a. 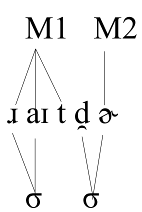

= /ɹaɪtɚ/ → [ɹaɪd̪ɚ]

b. 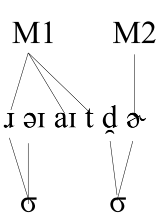

= /ɹaɪtɚ/ → [ɹəɪd̪ɚ]

M1 M2

c. *ɹ aɪ d̪ ɚ (= /ɹaɪɚ/ → [ɹaɪd̪ɚ])

σ σ

Because underlying /t/ will always be present in output representations, the constraint *[aɪ, aʊ]$_{pros}$[-voice] will always be violated whenever a non-raised diphthong precedes underlying /t/, regardless of whether this underlying /t/ is actually pronounced. This is shown in the tableau below. Because raising is motivated by non-pronounced voiceless segments, while flapping (triggered by *[VtV]$_{pros}$) only eliminates pronounced voiceless segments of a certain kind, the raising-triggering constraint and the flapping-triggering constraint can both be satisfied.

(98)

/ɹaɪtɚ/	*[aɪ, aʊ]$_{pros}$[-voice]	*[VtV]$_{pros}$	Faith
M1 M2 ɹ əɪ aɪ t d̪ ɚ = /ɹaɪtɚ/ → [ɹəɪd̪ɚ] σ σ			**
M1 M2 ɹ aɪ t d̪ ɚ = /ɹaɪtɚ/ → [ɹaɪd̪ɚ] σ σ	*!		*
M1 M2 ɹ əɪ aɪ t ɚ = /ɹaɪtɚ/ → [ɹəɪtɚ] σ σ		*!	*
M1 M2 ɹ aɪ t ɚ = /ɹaɪtɚ/ → [ɹaɪtɚ] σ σ	*!	*	

Whenever both *[aɪ, aʊ]$_{pros}$[-voice] and *[VtV]$_{pros}$ are ranked above relevant faithfulness constraints, the form with both raising and flapping will win, because it satisfies both markedness constraints. This could not be the case in OT with standard representations: because standard representations do not carry over all input material to outputs, but rather discard input material that has been changed or deleted, underlying unpronounced /t/ cannot be hypothesized to cause diphthong raising in Canadian English.

Thus, because representations in Colored Containment are richer than under standard representational assumptions, certain pieces of “contradictory” information (for instance, /t/ simultaneously with [d̬]) can actually both be present in output forms, enabling various constraints that evaluate outputs to refer to seemingly “contradictory” aspects of one and the same representation - which may give rise to opaque interactions (see the alternative definition of opacity in section 1.3).

4.1.4 Conceptual predecessors to differential classification

The current account of phonological opacity is built on the concept that sound events are independently classified at multiple taxonomic levels. In this section, I will mention selected examples of earlier results for which the presence of independent taxonomic levels of classifying sound events is crucial. These results come both from theoretical phonology and the experimental literature.

One almost trivial example of sound representations at different taxonomic levels co-existing with one another comes from the traditional divide between phonology proper and the phonetic component (Trubetzkoy 1939/1969, Kingston 2007). Both phonology and phonetics may be seen as having language-specific generalizations about sound shapes. However, the kind of units that are normally regarded as the units of generalization in both components differ. Phonetic generalizations are usually thought of as referring to continuous variables that map directly onto articulatory or acoustic dimensions. As the same time, phonological generalizations are often conceptualized as referring to variables with discrete levels, such as binary features (although see Browman & Goldstein 1986 and subsequent work in Articulatory Phonology for a dissenting perspective).

The discrete-level variables of phonology can be seen as being on a higher taxonomic level than the continuous variables of phonetics: a small number of phonological values summarizes an infinite number of phonetic values. Because the mapping from phonology to phonetics may not always be perfect, the phonetics-phonology interface is in some sense an implicit form of differential classification. In fact, some instances of opacity have been argued to arise from imperfect mapping between phonology and phonetics - see, for instance, Nazarov (2010).

Another body of work in which differences between taxonomic levels of classifying sound forms play a role is the theory of Lexical Phonology (Kiparsky 1982, 1985) - the intellectual precursor to Stratal OT. In this theory, phonology is roughly divided into a body of lexical rules and a body of post-lexical rules. Lexical rules apply within words, and are assumed to apply cyclically with the addition of every relevant morpheme. Post-lexical rules, however, apply after all morphology, and apply across words.

There are several intuitions about differences between phonological generalizations that have at various times been formalized in terms of the lexical/post-lexical distinction. One of these intuitions was formalized as the principle of Structure Preservation (Kiparsky 1985). This principle was hypothesized to hold of lexical, but not post-lexical rules, and its content can be roughly characterized as a ban on creating sounds that are not permitted in the lexicon of that language. In other words, lexical rules may only manipulate the contrastive sound units of a language, while post-lexical rules are free from this restriction.

When slightly re-conceptualized, this situation of Lexical Phonology with the Structure Preservation principle can be seen as follows: post-lexical rules manipulate universal phonological representations, while lexical rules manipulate language-specific phonological units. Language-specific phonological units classify universal phonological units in a particular way, and thus can be seen as a higher taxonomic level than universal phonological units. Since lexical rules and post-lexical rules, under this conceptualization, refer to the same sound events by different units - the lexical rules by higher level units and the post-lexical rules by lower-level units -, one can conclude that independent units at different taxonomic levels are indeed implicitly used in this framework.

There have also recently been some interesting experimental studies in speech perception and production which have found evidence for both lower taxonomic level and higher taxonomic level categories co-existing and co-determining listeners' and speakers' behavior. These studies include McQueen, Cutler & Norris (2006); Jesse, Page & Page (2007); and Nielsen (2011).

Here, I will briefly discuss Nielsen (2011), a study in which imitation of extraordinarily long (and extraordinarily short) aspiration in English stops was examined. It had been shown previously (Shockley, Sabadini & Fowler 2004, Pardo 2006) that exaggerated aspiration is imitated in the productions of subjects when they have been exposed to it. However, the interesting question posed by Nielsen is which level or levels of abstraction speakers use to generalize from the instances of exaggerated aspiration they have heard to new word tokens (produced by themselves).

Do speakers only generalize at the level of segment classes (such as the class of voiceless stops), so that exaggerated aspiration spreads equally to the production of all words starting with the same type of segment as the words heard, without regard for lexical identity or phoneme identity? Or do speakers only generalize at the level of phoneme categories, so that exaggerated aspiration is imitated only in words starting with the same phoneme as the words heard with exaggerated aspiration, but regardless of the lexical identity of the produced words? Or does generalization only take place at the level of individual words, so that words that have not been heard with exaggerated aspiration will not be pronounced with exaggerated aspiration?

Nielsen found that speakers' behavior did not provide evidence for only one of these levels of abstraction being used as the basis of generalization: instead, the degree of imitation of exaggerated aspiration was predicted by a combination of lexical, phonemic and featural factors.

The experiment relevant here (Nielsen's Experiment 1) was set up in three stages. First, participants were given a list of words (the “reading list”), and asked to read it out loud. Following this, they were exposed to a recording of another list of words in which every voiceless stop had exaggerated aspiration (the “listening list”). After exposure to the recording, participants were asked to read the reading list out loud once again. Comparing pre-exposure and post-exposure productions made it possible to see whether VOTs of voiceless stops were significantly lengthened after exposure to the listening list.

The listening list and the reading list matched and mismatched in such a way that the effects of lexical identity and frequency, and phonemic, and featural identity could be studied separately. The listening list only contained target words that started in /p/, half of which were high-frequency words, and the other half low-frequency words. The reading list contained all the target words from the listening list, plus a set of novel low-frequency words starting in either /p/ or /k/[25].

The results of this experiment showed that all target words on the reading list had significantly more aspiration after exposure to the recording. Thus, speakers must have generalized exaggerated aspiration to the class of voiceless stops. However, there was a separate effect for phoneme identity: words in the reading list that started in /p/ (i.e., in the same phoneme as the hyper-aspirated tokens on the listening list) exhibited significantly more increase in VOT than words that started in /k/. This same-phoneme advantage means that there must have been separate spreading of the exaggeration aspiration effect at the level of the phoneme /p/ only.

Note that the latter effect cannot be attained with featural representations only. It is true that /p/ shares all features with /p/, while it shares all but its place feature with /k/, which makes it appealing to explain the /p/-over-/k/ effect by assuming that the desire to imitate hyper-aspiration spreads along the dimensions of both the place feature and the manner features of /p/: that way, /p/-initial words receive more “aspiration activation” than /k/-initial words. However, Nielsen does not report that any sonorant-initial words in the reading list received aspiration – including /m/-initial words. If /m/ does not receive aspiration, how can exaggerated aspiration have spread along the feature [labial]? Thus, it appears that a separate level of individual phonemes is necessary to account for Nielsen's result.

Finally, there was an effect of lexical frequency: low-frequency words were hyperaspirated more often than high-frequency words. Even though there was no significant effect of lexical identity (words already encountered in the listening list did not show a larger increase in VOT than words not encountered in the listening list), the lexical frequency result still indicates that not all words are equally affected by exaggerated aspiration imitation, indicating that exaggerated aspiration is not only generalized at the level of phonemes and features, but is also mediated by some fine-grained level of episodic representation in the manner of exemplars (Goldinger 1998). Thus, the results indicate that featural, phonemic, and lexically-specific factors are involved in the imitation of exaggerated

25 Both lists also contained filler words that started in a sonorant.

aspiration.

According to Nielsen, “the crucial question to be asked is not whether phonological representations are abstract or episodic, but rather how abstract and/or episodic they are” (Nielsen 2011:133, author's emphasis). Her results establish that human processing of speech uses units at several levels of abstraction (at least episodic, phonemic, and featural) in parallel. Although this does not automatically entail that all these units are used in grammar (see, for instance, Chomsky 1965 on the dissociation of grammar and processing), the fact that different levels of abstraction from linguistic sounds must exist in the mind makes it much more plausible that different levels of abstraction may be used in phonological grammar.

4.2 Comparison with other theories: empirical predictions

The idea of differential classification is quite abstract in itself. However, it can still give rise to quite definite empirical predictions, both in the realm of phonological typology and in the realm of expected experimental findings. Typological predictions can be derived from differential classification if some kind of evaluation metric for phonological grammars is assumed. According to the original formulation of the evaluation metric for phonological grammars in Chomsky & Halle (1968), whenever two grammars describe the same primary linguistic data, only the grammar that has the shorter description may be chosen as the correct mental representation of these facts. I will extend this idea in the section below to apply to constraint-based grammars.

4.2.1 The evaluation metric and typological implications

Since differential classification introduces the concept that one and the same sound event will typically be represented by more than one separate unit (with segmental and with featural units), then constraints which license this sound event can be formulated in various ways: either referring to the segmental to the featural units that classify that sound event. For instance, when the sound events *[VpV] and *[VbV] are disallowed in a language, this can be expressed either by constraints banning the segmental units [p] and [b] in between two vowels, or by constraints that ban the feature bundle [+labial, -continuant, -sonorant] in between two vowels.

Assuming that at least some OT constraints are induced rather than innate (Hayes 1999, Hayes & Wilson 2008), it cannot be stipulated that the choice between a featural or a segmental formulation for a constraint that triggers a certain phenomenon is given by the innate grammar. This means that the language learner will need to find a way to establish for every phenomenon whether it is triggered by a segmental-level or a featural-level constraint.

One straightforward notion that seems useful in this context is the evaluation metric as espoused in early generative phonology (Chomsky & Halle 1968): when a phenomenon can be represented in the grammar in multiple ways, choose the grammatical representation with fewest symbols. Of course, it is not directly apparent how this could be translated into constraint-based terms. However, one way to formulate an evaluation metric for the situation outlined above (is phenomenon X triggered by a segmental-level or a feature-level constraint?) is the following:

Evaluation metric (for constraint induction)
Whenever a generalization requires inducing new constraints, choose the grammar which has fewer and simpler constraints given equal coverage of the generalizations.

The concept of preferring fewer and simpler constraints (while being somewhat under-formalized) can be illustrated by two contrasting cases: a case where a generalization targets one single segment versus a case in which a generalization refers to the entirety of a natural class of segments.

When a generalization states that only one particular segment (e.g., [t]) is prohibited or required in a particular context, this generalization can be expressed by one single constraint that refers only to that one segment category:*X[t]Y. Alternatively, this generalization could be expressed by a constraint or group of constraints that prohibit each of the features associated with that segment in the relevant context: *X[+coronal,-continuant,-sonorant,-voice,...]Y, or *X[+coronal]Y, *X[-continuant]Y, *X[-voice] etc.. This second, feature-level option would either require inducing a constraint that is more complex (refers to more representational units) than the segmental-level constraint, or inducing a larger number of constraints than would be the case when a segmental-level constraint were chosen. Thus, choosing a segmental-level representation of the ban on [t] in context X_Y would not require induction of as many constraints, or as complicated constraints as choosing a feature-level representation, which means that the Evaluation metric would require the ban on [t] in X_Y to be represented on the segmental level.

However, when there is a generalization that an entire natural class, for instance, the class of all sonorants, is banned in context X_Y, then the Evaluation metric would prefer the constraints that express this to be stated at the featural level. This can be shown by comparing potential segment-level and feature-level constraints against the configuration X-sonorant-Y. A segment-based formulation of such constraints would require enumerating all segments included in the class of sonorants, either in one complex constraint (referring to many disparate segmental units), or in a group of individual, simpler constraints: *X[m, n, ŋ, l, r, j, w, …]Y, or *X[m]Y, *X[n]Y, *X[l]Y, etc.. However, a feature-level formulation would only require one simple (in terms of the number of units referred to) constraint to express the ban of sonorants in the context X_Y: X[+sonorant]Y. For this reason, the Evaluation metric would only allow a feature-level formulation of this ban.

In this way, generalizations for which segment identity is more important than one or two particular featural dimensions will always be represented with constraints that refer to the segmental value of the target sound(s), while generalizations for which the converse is true (a very small number of particular features drives the generalization) will always be represented with feature-based constraints. This can be illustrated by the analysis of Yoruba as given in section 2 of this paper.

As has already been said at the end of section 2.1, it is plausible that the process full vowel assimilation in Yoruba is represented by constraints formulated in terms of segment units, because this process is only satisfied by full identity between adjacent vowels, not by partial feature correspondence between two vowels. At the other hand, ATR harmony in Yoruba is quite satisfied with partial identity between vowels, as long as this partial identity is on the tier of the feature [±ATR]; thus, the most satisfying grammatical representation of ATR harmony is in terms of features.

The Evaluation metric provides a reason to sharpen this intuition to a requirement: if the Evaluation metric holds, then Yoruba vowel assimilation will only be able to be represented by a segment-based constraint, while it will only be possible to represent ATR harmony in the same language with a feature-based constraints. Vowel assimilation, which requires that any two adjacent vowels be fully identical, can either be represented by one segment-based constraint against two non-identical segment values under certain circumstances (which is the constraint AssimV introduced in section 2.2)[26], or by a series of feature-based constraints, each of which requires identity of adjacent vowels on a certain feature dimension (*V[α back]V[-α back] “No adjacent vowels that do not agree in backness”, *V[α high]V[-α high] “No adjacent vowels that do not agree in being high vowels”, *V[α round]V[-α round], etc.), or by a single constraint which requires identity on all feature dimensions on which vowels can differ in Yoruba (height, backness, ATR, roundness, nasality): *V[α high, α round, α back, α nasal]V[-α high, -α round, -α back, -α nasal].

The latter two options require either having more than one constraint, or having a constraint which refers to many more representational units than AssimV does, so that these options are ruled out by the evaluation metric. In this manner, only grammars that contain a feature-based constraint like AssimV as the constraint triggering vowel assimilation are allowed to exist.

ATR harmony, as can be seen in section 2.1, requires that certain pairs of vowels in adjacent syllables be identical in their [±ATR] values, but not along other featural dimensions. For instance, the pairs [e … e], [e … o], [o … e], [o … o], [ɛ … ɛ], [ɛ … ɔ], [ɔ … ɛ], [ɔ … ɔ] are allowed because they agree in their [±ATR] value, but pairs such as *[e … ɛ], *[ɔ … e], *[ɛ … o] are not allowed.

To express this generalization on a segmental level, a host of constraints is needed - one against every prohibited combination of segments, of which there are 10 (as can be seen at the beginning of section 2) - *[e … ɛ], *[ɔ … e], *[ɛ … o], etc.. It is difficult to imagine how these 10 pairs could be fit into one complex constraint, but if this were to be done, then the resulting constraint would refer to a large number of individual phonological units, and would be considered complex.

26 Of course, the constraint AssimV is partially formulated in terms of features: the context in which non-identical segment values are to be avoided is given by the feature [+vocalic]; however, what is being compared in this case of application of the Evaluation metric is the formulation of the identity requirement between vowels, not the formulation of narrowing down the requirement to vowels only.

On the other hand, a featural interpretation of the facts allows for the formulation of just 2 constraints - *[-ATR][+ATR] and *[+ATR][-ATR], each of which refers to a small number of individual phonological units. It is easy to see that the Evaluation metric will prefer this variant of representing ATR harmony in Yoruba.

Thus, the Evaluation metric places rigid constraints on which phonological generalizations may be represented by segment-based constraints, and which may be represented by feature-based constraints. The account of phonological opacity developed in this paper requires that an opaque process A be represented in terms of a higher taxonomic level (such as features), while the process B that makes process A opaque must be represented in terms of a lower taxonomic level (such as segments). These two ideas combined entail that the possibility for two phonological generalizations to interact opaquely depends directly on their nature: only when the Evaluation metric allows one generalization to be represented by segment-based constraints, and another generalization to be represented by feature-based constraints, may these two generalizations interact opaquely.

This means that when the Evaluation metric requires process A to be represented by feature-based constraints, while process B must be represented by segment-based constraints, then process A cannot be made opaque by process B. An example of such a situation would be a fictional permutation of the Yoruba case, in which a feature-based process of vowel assimilation makes a segment-based vowel harmony process opaque. Let us call this fictional language Yoruba-prime.

Yoruba resolves the situation of adjacent non-identical vowels through full vowel assimilation. Instead of full vowel assimilation, Yoruba-prime has a process of vowel lowering, which turns the sequences /a+i/ and /a+u/ into [e] and [o], respectively. All other vowel-vowel sequences are unaffected.

(99) *hypothetical Yoruba-prime*
/abala ili/ → [abala eli]
/awaba uyu/ → [ayaba oyu]
/awaba ege/ → [awaba ege]

This process can be represented most compactly if it is triggered by a feature-based constraint *[+low][+high] (No directly adjacent elements of which the first is [+low] and the second is [+high].). If it were triggered by segmental-based constraints, then there would either have to be two constraints (*[a][u] and *[a][i]), or a single, but more complex constraint (*[a]{[i],[u]}: No sequences of [a] and then either [i] or [u].). The Evaluation metric will choose the feature-based interpretation.

At the same time, Yoruba-prime also has a vowel harmony process, like Yoruba does. However, vowel harmony in Yoruba-prime requires full identity of the vowels that are harmonized. Specifically, the second vowel of a word must be identical to the first vowel of that word (this is reminiscent of vowel harmony in some Romance dialects).

(100) *hypothetical Yoruba-prime*
/lo-ka/ → [loko]
/wi-lo-ka/ → [wilika]

If this process is to be triggered by feature-level constraints, then either a separate constraint would be necessary for each feature dimension that classifies vowels in this language, because the first two vowels of a word must be absolutely identical: Agree([±high]), Agree([±back], Agree([±ATR]), etc. Alternatively, one constraint could be created which would refer to all possible vowel features: Agree([±high, ±back, ±ATR, …....]).

A segment-based formulation of the constraint triggering this harmony process, however, would be very simple: a single constraint of the shape Agree(segment)[27] would suffice. Because this is a single constraint which only refers to one instead of many tiers of representation, the Evaluation metric will only allow this segment-based formulation of the constraint triggering vowel harmony.

In analogy to Yoruba, Yoruba-prime has an opaque interaction between the two processes laid out above: vowel lowering makes vowel harmony opaque in this hypothetical language. When two vowels at the edge of a word are coalesced, the resulting vowel does not trigger vowel harmony.

(101) *hypothetical Yoruba-prime*
/abala ili/ → [abala eli], *[abala ele]
/awaba uyu/ → [awaba oyu], *[awaba oyo]

However, this opaque interaction cannot be represented under a differential classification approach. The tableau below shows the two constraints which trigger vowel coalescence and vowel harmony, respectively: *[+low][+high] and Agree(segment), as well as the feature-to-segment association constraint [e/o] → [-high]. Faithfulness to segments must be dominated by Agree(segment), since the harmony requirement may change underlying segment values (as shown in (100) above).

Even when [e/o] → [-high] is given the lowest possible ranking, and *[+low][+high] is given the highest possible ranking to indicate that coalescence is more important than harmony, there is no way to derive the opaque interaction shown above: the opaque mapping /abala ili/ → [abala eli] cannot win from the transparent mapping /abala ili/ → [abala ele], unless Faith(segment) is ranked above Agree(segment) - but that would mean that there is no harmony at all in the language, which is false.

27 Of course, this sidesteps the issue of restriction of harmony to the first two syllables, and the directionality of harmony.

(102)

[+lo] [+hi] [+hi] /a ba la i li /	*[+low][+high]	Agree(segment)	Faith(segment)	[e/o] → [-high]
[+lo] [-hi] [-hi] a ba la e le			**	
[+lo] [-hi] [+hi] a ba la e li		*!	*	
[+lo] [+hi] [+hi] a ba la e li		*!	*	*
[+lo][-hi] [+hi] a ba la i li	*!			

In this tableau, the transparent mapping /abala ili/ → [abala ele] wins because the opaque candidates (candidates 2 and 3) both violate Agree(segment), and the faithful candidate /abala ili/ → [abala ili] violates *[+low][+high]. The fact that the third candidate has identical feature values for the first two vowels in the word [eli] does not prevent this candidate from violating Agree(segment), since this constraint only looks at the segment tier, and the identical feature values on the two vowels do not improve anything. Compare this tableau to the tableaux at the end of section 2, where a mismatch between features and segments.

In conclusion, because an opaque interaction such as the one above would involve a segmental process made opaque by a featural process, Yoruba-prime is predicted to be an impossible language by the differential classification account of opacity. However, other accounts of opacity would easily permit such a language. Stratal OT would permit this interaction, because the opaque word-bounded vowel harmony process operates over a smaller morphosyntactic domain than vowel lowering, which operates across words. There is also no reason why this interaction could not be analyzed in OT-CC or Containment-based approaches.

Differential classification hypothesizes that opaque interactions between phonological generalizations can only occur when the opaque generalization refers to a higher taxonomic level of phonological representation than the generalization which makes opaque. Given the Evaluation metric as a diagnostic tool to determine the appropriate taxonomic level for a generalization, differential classification imposes definite constraints on the typology of opaque interactions: a process which targets specific segments cannot make a process which targets featural dimensions opaque.

It remains to be investigated whether this restriction is a good fit to actual language typology. However, at the very least the theory of differential classification brings forward a clear empirical hypothesis which can be refuted by linguistic fact.

4.2.2 Implications for experimental research

Beside the typological claim made by differential classification, as sketched above, the approach of differential classification also generates expectations for outcomes of experimental studies. As has been shortly summarized in section 4.1.4 above, experiments in speech perception and production have been able to find evidence for different taxonomic levels in the processing of speech. The model of opacity developed here builds on this concept of taxonomic levels: opaque generalizations will refer to a higher taxonomic level than the generalizations that make them opaque, and the representation of the elements crucial to the opaque interaction will have an "inconsistency" between the lower and the higher taxonomic level. For instance, vowels whose quality is adjusted by vowel assimilation in Yoruba will have "inappropriate" feature values.

This setup predicts that experiments in which it is investigated along which dimensions and to which extent a certain speech idiosyncrasy is extended to new items by listeners (cf. the Nielsen 2011 study cited in section 4.1.4) will find that speech idiosyncrasies presented on segments crucial to opaque interactions will spread to new segments that are in a "mismatching" natural class. For instance, when the Yoruba vowels which undergo vowel assimilation such as the first [o] in /εba odo/ → [εbo̲ odo] are presented to listeners with some idiosyncrasy (e.g., harsh voice), it is predicted that this idiosyncrasy will spread more strongly to vowels that have the same underlying feature values (in this case, to other low vowels) rather than to other segments that have a feature value typically appropriate for that segment (in this case, to other mid vowels, or to other round vowels).

(103)

[+lo]	[+lo]	
/ε ba̲ odo/ →	[ε bo̲ odo]	Prediction: harsh voice on first [o] will be generalized to low vowels more than to mid vowels.

However, this behavior is only predicted for segments that must be differentially classified. When a segment undergoes change, but there is no reason to assume that there is a mismatch between segmental and feature value, it is predicted that an idiosyncrasy presented on that segment will not spread to segments that share the underlying features of that segment, but rather to segments which share the expected surface feature values of that segment. For instance, when ATR harmony changes a vowel, it is expected (given the analysis in section 2) that the underlying [±ATR] value of that vowel will change. When an idiosyncrasy is introduced on a vowel of this sort, it will be predicted to spread more strongly to new tokens of vowels that share the surface-predicted features of that vowel.

(104)
[+ATR] [-ATR]   [-ATR] [-ATR]
/ e rɔ / → [ɛ rɔ] Prediction: harsh voice on [ɛ] will be generalized to [-ATR] vowels more than to [+ATR] vowels.

Thus, differential classification predicts that segments that must be differentially classified in the theory will behave differently from segments that are thought to have standard feature-segment associations. This difference in behavior is predicted to be observed in any experimental paradigm which probes different taxonomic levels of classifying sound events.

4.2.3 Comparison to predictions of other theories

When the empirical predictions of differential classification, as briefly sketched above, are considered, it becomes apparent that there is no evident way that any of the three other approaches to opacity introduced above (Stratal OT, OT-CC, and Containment) would lead to the same empirical predictions. While it is unclear whether the empirical predictions sketched above are to any extent true of natural language, they do clearly distinguish differential classification from other theories of opacity. In what follows, I will discuss some of the empirical predictions made by each of the three other approaches (Stratal OT, OT-CC and Containment), and show what differential classification has to say with respect to these predictions.

As was discussed in section 4.1.1, Stratal OT hypothesizes that opaque interactions may arise when two phonological generalizations are represented in different strata of phonological grammar, and the opaque generalization is at a stratum which applies to a smaller morphosyntactic domain than the generalization which makes opaque. Of course, this notion is very abstract: the concept that a phonological generalization holds over inputs of a certain morphosyntactic size does not correspond one-to-one with the actual morphosyntactic domain of the generalization.

For instance, a word-bounded generalization may be stated in the stratum of phrase-level phonology: because phrases contain words, the constraints triggering this generalization may refer to word-size domains within phrases. Speaking more generally, any process which is bounded by a certain morphosyntactic domain may be expressed at a phonological stratum that processes inputs of a larger morphosyntactic size (stem-bounded generalizations may be expressed in word-level or phrase-level phonology).

However, it is not the case that a generalization which applies across words can be represented in word-level phonology and not in phrase-level phonology. In general, a process that applies to a larger morphosyntactic unit cannot be represented only in a stratum which processes inputs of a smaller morphosyntactic size. Thus, processes that apply across words cannot be represented in word-level or stem-level phonology only. Similarly, processes that apply across any morpheme boundary in words may not be represented in stem-level phonology.

From this consideration, and the fact that Stratal OT requires that the opaque generalization in an opaque interaction be at a stratum of lower morphosyntactic size than the generalization which makes opaque, it follows that word-bounded or stem-bounded processes may not make a process that applies across words opaque. For instance, if Canadian Raising applied across words, but flapping was word-bounded, Stratal OT would have no way of representing the opaque interaction between the two processes.

There are actually some cases of opacity that are of this kind. For instance, McCarthy (2007b) shows that Bedouin Arabic has an interaction (not analyzed in section 3 of this paper) between vowel raising and high vowel syncope which contradicts the predictions of Stratal OT. Vowel raising makes high vowel syncope opaque; however, vowel raising is a word-bounded process, while syncope operates across words. For this, it follows that syncope could not be represented at a stratum that is lower than phrase-level. However, for syncope to be made opaque by vowel raising, vowel raising needs to be represented in a stratum that applies to a larger morphosyntactic domain - but this is not possible, since it is assumed in Stratal OT that the phrase is the largest morphosyntactic domain that has a separate stratum of phonological grammar associated with it.

Thus, it appears that the prediction made by Stratal OT that non-word-bounded processes should not be opaque is not borne out. This is an advantage for differential classification, since differential classification does in any way connect morphosyntactic domains to phonological opacity: processes that are in opaque interaction may be associated with any morphosyntactic domain, as long as the opaque process refers to a higher taxonomic level than the process which makes opaque.

The empirical predictions with regard to the typology of opaque interactions made by OT with Candidate Chains are complex, but I will highlight two predictions. One of these is that opaque generalizations should never be static phonotactic generalization, but rather should involve some fixed process of change. The other prediction is that so-called “counterfeeding from the past” (Wilson 2006, Wolf 2010) should be possible. Both of these predictions will be explained below, and compared to the predictions of differential classification.

The mechanism which drives opacity in OT-CC is the concept of ordering unfaithful mappings of different kinds: a generalization is opaque because the change it is associated with may not take place after some other kind of process has applied. If a phonological generalization has no process of change associated with it - in other words, when a generalization is a static phonotactic generalization - there is no way to express the fact that this generalization is enforced “before” some other generalization: no Precedence constraint could demand this, since Precedence constraints are formulated in terms of unfaithful mappings.

Thus, OT-CC predicts that there would never be static phonotactic regularities that are made opaque by some process. Of course, this claim is true to the extent that it is truly impossible to identify the operation of change that would be used to enforce a phonotactic regularity, or when the typical operations used to enforce this regularity cannot be coherently ordered before the operation belonging to the process that makes opaque.

It turns out that static phonotactic regularities can indeed be made opaque. For instance, in Aymara (Adelaar & Muysken 2004), every stem must end in a vowel, but there is a morphosyntactically conditioned process of vowel deletion which deletes vowels, disrespecting all phonotactic restrictions.

(105)
kunturi "condor" (< Proto-Quechua *kuntur)
hawasa "beans" (< Spanish habas)
kumpari "compadre" (< Spanish compadre)
umalsu "(place name)"

(106)
/č'iyara uta/ → č'iyar uta "black house"
/čura-ta-χa/ → čur-ta-χa "you gave it" (give-2sg-topic)

If both of these regularities are given a place in phonological grammar, then it follows that the final vowels in (106) are present for the sake of phonotactics (morphotactics) while they are absent for vowel deletion. This creates an opaque interaction. At the same time, there is no particular process of change that the requirement of vowel-finality is associated with. This makes it at least apparently impossible to represent this opaque interaction using a Precedence constraint.

Regardless of whether such opaque phonotactic regularities, upon further scrutiny, could be reanalyzed using Precedence constraints, it should be noted that differential classification makes no distinction between phonotactic regularities which do not have a fixed "repair strategy" associated with them, or alternation-triggering phonological processes which do have such a fixed strategy of change. All that matters for differential classification is that one generalization, static or dynamic, be at a different taxonomic level than another generalization, static or dynamic.

Another prediction made by OT-CC is that some opaque interactions will have the shape of so-called "counterfeeding from the past" (Wilson 2006, Wolf 2010). Such interactions are predicted to arise from interactions of Precedence constraints, and cannot be described in terms of rule ordering: process B may not feed process C if some process A, which must be ordered before B, has applied (see Wolf 2010 for examples of such interactions in natural language). Such interactions cannot be described in the theory of differential classification. However, the number of cases that are attested to be of this type is very small, and it would be worth examining these cases and seeing if there is some additional factor (e.g., prosody) which could have caused this "counterfeeding from the past" behavior.

Finally, Containment-based approaches to opacity also make certain predictions. One prediction made by such theories in general is that prosodic generalizations will never be opaque.

The latter follows from the premise that prosodic parsing determines the surface form in Containment-style phonology. This is also false: stress placement may be maybe opaque by some other deletion or insertion operation (see Elfner 2009, to appear for Dakota). Differential classification, at this point of development, remains agnostic about opaque prosody, as such cases have not yet been systematically investigated in this framework - this is left for future research.

References

Adelaar, W. & P.C. Muysken. 2004. *The Languages of the Andes*. Cambridge: Cambridge University Press.

Al-Mozainy, H.Q. 1981. *Vowel Alternations in a Bedouin Hijazi Arabic Dialect: Abstractness and Stress*. Doctoral dissertation, University of Texas, Austin.

Archangeli, D. & D. Pulleyblank. 1989. "Yoruba vowel harmony." *Linguistics Inquiry*, 20, 2, 173-217.

Baković, E. 2003. "Vowel harmony and stem identity." *San Diego Linguistic Papers, 1*. San Diego (CA): Department of Linguistics, UCSD.

Baković, E. 2011. "Opacity and ordering." In: Goldsmith, J., J. Riggle & A.C.L. Yu (eds.), *The handbook of phonological theory: second edition*, Oxford: Blackwell Publishing, pp. 40-67.

Bamgbose, A. 1966. *A Grammar of Yoruba*. Cambridge: Cambridge University Press.

Bermúdez-Otero, R. 2004. "Raising and Flapping in Canadian English: grammar and acquisition." Handout of paper presented at the CASTL Colloquium, University of Tromsø, 2 November 2004. [Available at www.bermudez-otero.com/tromsoe.pdf]

Bermúdez-Otero, R. 1999. *Constraint interaction in language change: quantity in English and Germanic* [*Opacity and globality in phonological change*]. PhD dissertation: University of Manchester & Universidad de Santiago de Compostela.

Bermúdez-Otero, R. 2003. "The acquisition of phonological opacity." In: J. Spenader, A. Eriksson & Ö. Dahl (eds.), *Variation within Optimality Theory: Proceedings of the Stockholm Workshop on 'Variation within Optimality Theory'*, Stockholm: Department of Linguistics, Stockholm University, pp. 25-36.

Browman, C. & L. Goldstein. 1986. "Towards an articulatory phonology". *Phonology Yearbook*, 3, 219-252.

Chomsky, N. 1964. "Current issues in linguistic theory. In: J. Fodor & J. Katz (eds.), *The structure of language*, Eaglewood Cliffs (NJ): Prentice Hall, pp. 50-118.

Chomsky, N. 1965. *Aspects of the theory of syntax*. Cambridge (MA): M.I.T. Press.

Chomsky, N. & M. Halle. 1968. *The sound pattern of English*. New York (NY): Harper and Row.

Elfner, E. 2009. “Syllabification and stress-epenthesis interactions in Harmonic Serialism”. Ms., University of Massachusetts, Amherst. Available on Rutgers Optimality Archive, #1047.

Elfner, E. to appear. “Stress-epenthesis interactions in Harmonic Serialism”. To appear in: J. McCarthy & J. Pater (eds.), *Harmonic Grammar and Harmonic Serialism,* London: Equinox Press.

Gafos, A. & L. Lombardi. 1999. “Consonant transparency and vowel echo.” In: P. Tamanji, M. Hirotani, & N. Hall (eds.), *Proceedings of the Northeast Linguistic Society 29*, Amherst (MA): Graduate Linguistics Student Association, University of Massachusetts, pp. 81-95.

Gouskova, M. & Hall, N. 2009. “Acoustics of epenthetic vowels in Lebanese Arabic.” In: S. Parker (ed.), *Phonological argumentation: Essays on evidence and motivation,* London: Equinox Publishing, pp. 203–225.

Hayes, B. 1999. “Phonetically-Driven Phonology: The Role of Optimality Theory and Inductive Grounding”. In: M. Darnell, E. Moravscik, M. Noonan, F. Newmeyer, and K. Wheatly (eds.), *Functionalism and Formalism in Linguistics*, Volume I: General Papers, Amsterdam: John Benjamins, pp. 243-285.

Hayes, B. & C. Wilson 2008. “A maximum entropy model of phonotactics and phonotactic learning”. *Linguistic Inquiry*, 39, 379-440.

Jesse, A., J.M. Page & M. Page. 2007. “The locus of talker-specific effects in spoken-word recognition.” In: *Proceedings of ICPhS XVI*, pp. 1921-1924.

Joos, M. 1942. “A phonological dilemma in Canadian English.” *Language*, 18, 141-144. Reprinted in: J.K. Chambers (1975), *Canadian English: Origins and Structures*, Toronto: Methuen, pp. 79–82.

Kingston, J. 2007. “The phonetics-phonology interface.” In P. de Lacy (ed.), *The Cambridge Handbook of Phonology,* Cambridge: Cambridge University Press, pp.435-456.

Kiparsky, P. 1973. “Abstractness, opacity, and global rules.” In: O. Fujimura (ed.), *Three Dimensions of Linguistic Theory*, Tokyo: TEC, 57–86.

Kiparsky, P. 1982. “From Cyclic Phonology to Lexical Phonology.” H. v. d. Hulst & N. Smith (eds.), *The Structure of Phonological Representations*, Dordrecht.

Kiparsky, P. 1985. “Some Consequences of Lexical Phonology.” *Phonology Yearbook*, 2, 85-138.

Kiparsky, P. 2000. “Opacity and Cyclicity.” *The Linguistic Review*, 17, 351-367.

Lombardi, L. 1995. *Why place and voice are different: Constraint interactions and featural faithfulness in Optimality Theory. Ms., University of Maryland, College Park.*

McCarthy, J. 2007a. “Derivations and levels of Representation.” In: P. de Lacy (ed.), *The Cambridge Handbook of Phonology*, Cambridge: Cambridge University Press, pp. 99-117.

McCarthy, J.J. 1999. “Sympathy and phonological opacity.” *Phonology*, 16, 331-399.

McCarthy, J.J. 2007b. *Hidden Generalizations: Phonological Opacity in Optimality Theory*. London: Equinox.

McCarthy, J. & A. Prince. 1993. *Generalized Alignment*. Unpublished ms., University of Massachusetts, Amherst, and Rutgers University, New Brunswick.

McCarthy, J. & A. Prince. 1994. “The emergence of the unmarked: Optimality in Prosodic Morphology.” In: M. Gonzàlez (ed.), *Proceedings of the North East Linguistics Society 24*, Amherst (MA): Graduate Linguistics Student Association, University of Massachusetts, pp. 333-379.

McCarthy, J.J. & A. Prince. 1995. “Faithfulness and reduplicative identity.” In: J. Beckman, L.W. Dickey & S. Urbanczyk (eds.), *Papers in Optimality Theory*, University of Massachusetts Occasional Papers 18, Amherst (MA): Graduate Linguistics Student Association, pp. 249–384.

McQueen, J.M., A. Cutler & D. Norris. 2006. “Phonological abstraction in the mental lexicon.” *Cognitive Science*, 30, 1113-1126.

Mielke, J. 2004. *The emergence of distinctive features*. Doctoral dissertation, Ohio State University.

Nazarov, A.I. 2010. *Phonological opacity emerges from Strictly Parallel Optimality Theory*. M.Phil. thesis, Leiden University.

Nevins, A. 2004. *Conditions on (dis)harmony*. Doctoral dissertation, Massachusetts Institute of Technology.

Nielsen, K. 2011. “Specificity and abstractness in VOT imitation”. *Journal of Phonetics, 39, 132-142*.

Oostendorp, M. van. 2004. *The theory of faithfulness*. Ms., Meertens Instituut.

Oostendorp, M. van. 2007. “Derived Environment Effects and Consistency of Exponence.” In: S.

Blaho, P. Bye & M. Krämer (eds.), *Freedom of Analysis?*, Berlin: Mouton de Gruyter, pp. 123-148.

Oostendorp, M. van. 2008. “Incomplete Devoicing in Formal Phonology.” *Lingua*, 118, 1362-1374.

Pardo, J. 2006. “On phonetic convergence during conversational interaction”. *Journal of the Acoustical Society of America*, 119, 2382–2393.

Pater, J. to appear. “Canadian raising with language-specific constraints.” *Language*.

Piggott, G. 1995. “Epenthesis and syllable weight.” *Natural Language and Linguistic Theory*, 13, 283-326.

Prince, A. & P. Smolensky. 1993/2004. *Optimality Theory: Constraint interaction in generative grammar.* Technical report, Rutgers University and University of Colorado at Boulder, 1993. ROA 537, 2002. Revised version published by Blackwell, 2004.

Pulleyblank, D. 1988. “Vocalic underspecification in Yoruba.” *Linguistic Inquiry*, 19, 2, 233-270.

Revithiadou, A. 2007. “Colored Turbid accents and Containment: A case study from lexical stress”. In: Blaho, S., P. Bye & M. Krämer (eds.), *Freedom of Analysis?*, Berlin & New York: Mouton de Gruyter, pp. 149-174.

Shockley, K., Sabadini, L., & Fowler, C.A. 2004. “Imitation in shadowing words”. *Perception & Psychophysics*, 66, 3, 422–429.

Smolensky 1996. *The Initial State and “Richness of the Base”*. Technical Report JHU-CogSci-96-4, Department of Cognitive Science, John Hopkins University.

Trubetzkoy, N.S. 1939/1969. *Principles of Phonology*. (Translation of *Grundzüge der Phonologie*.) Berkeley (CA): University of California Press.

Wilson, C. 2006. *Counterfeeding from the past*. Ms., University of California, Los Angeles.

Wolf, M. 2010. “On the existence of counterfeeding from the past”. Talk given at the 84th LSA Annual Meeting, Baltimore (MD), January 9, 2010.