Dissertation Proposal

Eileen Blum[1]

1 Introduction

I propose to investigate the computational complexity of vowel harmony patterns using forbidden substructure constraints (FSCs) over multi-tiered autosegmental representations (ARs). This investigation provides a well-defined, computationally motivated theory of well-formedness in vowel harmony. This follows Jardine (2017) who developed a theory of tonal well-formedness and determined that tone patterns are fundamentally local over two-tiered ARs. Investigating the locality of vowel harmony patterns in this way allows for a restrictive theory of well-formedness that aims to sufficiently capture the attested typology.

I analyze vowel harmony as a phonotactic restriction using only surface ARs rather than an input-output map. Previous work has analyzed vowel harmony patterns as resulting from a single assimilation process, whether it be feature spreading or agreement (Bakovic, 2000; Clements, 1976; McCarthy, 2011; Nevins, 2010; Rose & Walker, 2011; van der Hulst & Smith, 1986; Walker, 2010). Like Walker (2011), I show that given a local theory of surface markedness constraints, we expect vowel harmony patterns to utilize surface ARs that reflect either type of assimilation. I propose to analyze a wide range of vowel harmony patterns to determine whether or not they can be captured by FSCs, which offer a uniform theory of surface markedness based in a well-defined notion of locality.

I present two case studies of vowel harmony patterns with neutral vowels. Blocking in Akan is easily captured by a single FSC, which forbids an AR with two different ATR features on one tier and only a single +low feature on another tier. Transparent vowels in Finnish are captured—without relying on language-specific featural underspecification—by FSCs that forbid an AR with two different back features when the [-back] vowel is also associated to either [+low] or [+round]. These two cases demonstrate the locality of vowel harmony over multi-tiered ARs. In addition, these cases provide examples of the two assimilation mechanisms I claim to be present in vowel harmony.

The theory of vowel harmony I propose differs from Walker (2011)'s in that it does not rely on correspondence. Walker (2011) proposes a theory of vowel harmony with transparent vowels, which relies on correspondence between non-adjacent identical features on a tier. However, the ARs used in the Finnish case study illustrate that—on the surface—harmony propagation is demonstrated between non-adjacent features on a tier because they have the same + or - value, which differs from the intervening feature value on that same tier; and there is no need to connect identical features with an additional correspondence relation. In order to compare my theory to Walker (2011)'s theory, I will analyze a wide range of patterns that she captures to determine the explanatory adequacy of my theory. I highlight two patterns that provide important insights for phonological theory: first-last harmony in Eastern Meadow Mari and length-conditioned harmony in Baiyina Orochen.

To understand the typological predictions of my theory, I aim to situate vowel harmony within Heinz (2018)'s subregular hierarchy. Heinz (2018) proposes that all attested vowel harmony patterns are Strictly Local (SL) or Strictly Piecewise (SP), but Aksënova (2017) argues that vowel harmony patterns with transparent vowels are Tier-based Strictly Local (TSL) over strings. Heinz (2018) further claims first-last harmony is unattested because it is Locally Testable (LT) over strings and his SL/SP theory predicts LT patterns to be unattested. However, Walker (2011) describes a vowel harmony pattern in Eastern Meadow Mari that behaves like first-last harmony. In addition, enriching the representation has been shown to decrease the complexity of a pattern. I propose to analyze these and other vowel harmony patterns over multi-tiered ARs in order to determine the relationship between patterns describable with FSCs over multi-tiered ARs and existing complexity classes.

A reach goal is to investigate the newly developed quantifier-free least-fixed point logic to investigate how local surface phonotactic constraints over muti-tiered ARs inform a transformational analysis of vowel harmony patterns (Chandlee & Jardine, 2019; Koser, Oakden, & Jardine, 2019).

1.1 The Subregular Hierarchy

A goal shared by all of generative phonology is to distinguish attested patterns from logically possible, but unattested ones. A theory of well-formedness in vowel harmony that accomplishes this goal must be both expressive enough to explain the attested typology of vowel harmony patterns and restrictive enough to exclude the logically possible unattested vowel harmony patterns. This goal of distinguishing attested phonological patterns from possible unattested patterns is currently being investigated using formal language theory to determine the expressive power required to compute phonological patterns in general. A significant body of work in computational phonology shows that phonological generalizations are properly contained within the regular class of stringsets (Heinz & Idsardi, 2013). Recent work has further established a subregular hierarchy of stringset classes, i.e. star-free (SF) and weaker classes (Heinz, Rawal, & Tanner, 2011; Rogers & Pullum, 2011; Rogers et al., 2013). The subregular hierarchy, in (1), classifies stringsets in terms of the relative expressivity of the grammars needed to generate them. Each class that is lower on the hierarchy is also a proper subset of the class that dominates it, connected by a straight line.

(1) The Subregular Hierarchy (adapted from Heinz, 2018)

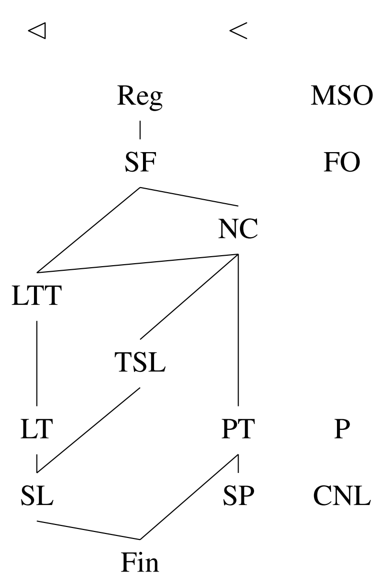

Stringset classes: Regular (Reg), Star-Free (SF), Non-Counting (NC), Locally Threshold Testable (LTT), Locally Testable (LT), Piecewise Testable (PT), Strictly Local (SL), Tier-based Strictly Local (TSL), Strictly Piecwise (SP), Finite (Fin)
Ordering Relations: Successor ($\lhd$), Precedence ($<$)
Logical Power: Monadic Second Order (MSO), First Order (FO), Propositional (P), Conjunctions of Negative Literals (CNL)

The diagram in (1) also illustrates the correlation between the expressive power of a logic and the subregular class of patterns that logic can describe. On the right side of the diagram are the minimal levels of logic needed to describe patterns in each stringset class. For example, Heinz (2018) claims that phonotactic constraints are either SL or SP—depending on which ordering relation is used—and they can be described using conjunctions of negative literals. I will adopt the formal language theory approach, which provides explicit ways of determining the locality of vowel harmony patterns. I will also use this approach to investigate whether the surface well-formedness theory I propose must be restricted to avoid predicting unattested patterns.

There is a type of phonotactic constraint, which picks out a piece of a string that is forbidden (or marked) and whose presence will make a string ungrammatical. Phonologists often write such constraints as OT markedness constraints with an asterisk (*), but these markedness constraints can be easily translated into the lowest level of logic, conjunctions of negative literals (CNLs) written with Boolean logical connectives, such as $\neg$ and $\wedge$. For example, in (2) a stringset ($\mathbb{L}_1$) contains only strings made up of the alphabet (Σ) of two symbols, *a* and *b*, and these symbols are connected by the successor ordering relation ($\triangleleft$). The stringset $\mathbb{L}_1$ can be characterized by the phonotactic constraints $^*aa$ and $^*bb$. These two constraints can be translated into a single CNL, which reads as "Strings which do not contain *aa* as a substring and do not contain *bb* as a substring are well formed"(Heinz, 2018). This stringset, $\mathbb{L}_1$, consists only of strings with either a single *a*, a single *b*, or alternating *a*s and *b*s. It contains no strings with the substring *aa* or *bb*, which means that two of the same symbol can not be adjacent within a string. Another stringset, $\mathbb{L}_2$, can be characterized by one of those constraints, $^*bb$, interpreted over strings of *a*s and *b*s connected with the precedence relation (<). The CNL for $\mathbb{L}_2$ will read as "Strings which do not contain *bb* as a subsequence are well formed"(Heinz, 2018). The stringset $\mathbb{L}_2$ in (3) thus consists only of strings with a single *b* because it contains no strings with the subsequence *bb*; two *b*s cannot occur at any distance within the string.

(2) SL_2 phonotactic constraints over strings
$\langle \Sigma = \{a,b\}, \triangleleft \rangle$ CNL = $\neg aa \wedge \neg bb$
$\mathbb{L}_1(\neg aa \wedge \neg bb) = \{a, b, ab, ba, aba, bab, abab, baba, ...\}$

(3) SP_2 phonotactic constraints over strings
$\langle \Sigma = \{a,b\}, < \rangle$ CNL = $\neg bb$
$\mathbb{L}_2(\neg bb) = \{a, aa, b, ab, ba, aaa, aab, baa, ...\}$

Unlike the basic phonotactic restrictions modeled in (2) and (3), Heinz (2018) argues that certain logically possible phonological patterns are unattested because they must be categorized in a higher subregular stringset class. For example, he categorizes a first-last harmony pattern evaluated over strings in the Locally Testable (LT) class, which requires the more expressive propositional logic. There are straightforward computational ways to determine the class into which a stringset falls and (4) utilizes diagnostics from Rogers et al. (2013) to demonstrate why a first-last harmony pattern over strings is neither SL nor SP.

(4) First-Last Harmony (FLH)
   a. not SL: $ab^{k-1}a \in$FLH and $cb^{k-1}c \in$FLH, but $ab^{k-1}c \notin$FLH
   b. not SP: $ab^m c^n a \in$FLH and b...c$\sqsubseteq ab^m c^n a$, but b...c$\notin$FLH

Heinz (2018) claims that first-last harmony patterns over strings are unattested. A theory of phonotactics that claims all cooccurence restrictions can be described using CNLs and are thus SL or SP explains that first-last harmony patterns would be unattested because they are neither SL nor SP over strings for any k, as shown in (4). To support this theory, Lai (2015) provides experimental evidence of a learning bias toward learning an attested SP pattern rather than the supposedly unattested LT first-last pattern.

1.2 Tiers and non-local patterns

Heinz et al. (2011) argues for a tier-based class of patterns to account for patterns like vowel harmony and Aksënova (2017) argues that long-distance vowel harmony patterns are TSL (Aksënova & Deshmukh, 2018). A TSL grammar is one in which a subset of elements in a string are projected onto a separate tier. The elements on a tier projection can be a subset of the elements in the original string or different tiers can have disjoint sets of elements, but no tiers consist of intersecting sets of elements. Under a TSL analysis vowel harmony is evaluated on a tier that includes only a subset of the elements in the original string. Aksënova (2017) claims that ATR harmony in Lokaa with transparent high vowels is thus local on a tier of non-high vowels. The generalization for Lokaa is that “a non-high vowel agrees with the preceding non-high vowel in ATR” and “all high vowels and consonants are transparent for the harmony” (Akinlabi, 2009; Aksënova, 2017; Aksënova & Deshmukh, 2018).

(5) Lokaa (Niger-Congo)
  a. èsìsòn ‘smoke’ *èsìsɔ̀n
  b. ɛ̀sísɔ̀n ‘housefly’ *ɛ̀sísòn
  c. lèjìmə̀ ‘matriclan’ *lèjìmà
  d. ɛ́kílìkà ’kind of plant‘ *ékílìkà

(6) ATR harmony in Lokaa

e				o		*	e				ɔ		
|				|			|				|	
è	s	ì	s	ò	n		è	s	ì	s	ɔ̀	n	

In order for ATR harmony to skip over the transparent high vowels and consonants, a separate tier is projected, which contains only the non-high vowels, as in (6). Then the Lokaa ATR harmony pattern can be described using a CNL that forbids two adjacent elements with different ATR feature specifications and that constraint is evaluated only over the tier projection of non-high vowels. The TSL grammar that generates the Lokaa vowel harmony pattern would look like (7).

(7) TSL grammar for Lokaa harmony

  a. Tier of non-high vowels: T = ɛ, e, o, ə, ɔ, a
  b. CNL: $\neg[\alpha \text{ATR}][\beta \text{ATR}]$ = ¬ɛe ∧ ¬eɛ ∧ ¬ɛo ∧ ¬oɛ ∧ ¬ɛə ∧ ¬əɛ ∧ ¬ɔe ∧ ¬eɔ ∧ ¬ɔə ∧ ¬əɔ ∧ ¬ɔo ∧ ¬oɔ ∧ ¬aə ∧ ¬əa ∧ ¬ao ∧ ¬oa ∧ ¬ae ∧ ¬ea

One other aspect of the TSL grammar in (7) is that elements in a string which do not participate in or affect harmony–so-called transparent elements, such as high vowels and consonants in Lokaa—are excluded from the tier. This allows the pattern to be described with FSCs over a subset of the segments in a word.

So, the computational complexity of a pattern depends upon the set of representational primitives over which a pattern is evaluated. Utilizing a more complex representation reduces the computational power and logical expressivity needed to evaluate a pattern. Tone patterns, for example, become ASL when evaluated over ARs rather than strings (Jardine, 2016, 2017). Some previous analyses of vowel harmony also use ARs to represent the spreading of a vowel feature from one vowel throughout the word until it is blocked (Clements, 1976; Goldsmith, 1976; McCarthy, 1988; Padgett, 2002; Sagey, 1986; van der Hulst, 2017; Walker, 2010, 2011, 2014). I propose to investigate the complexity of vowel harmony patterns over multi-tiered ARs rather than assuming all vowel harmony patterns are TSL.

There are two important motivations for investigating the effect of multi-tiered ARs on the computational complexity of vowel harmony patterns. First, Heinz (2018) claimed that first-last harmony patterns are unattested, but Walker (2011) analyzes Eastern Meadow Mari, which is described as utilizing a first-last harmony pattern as well as a number of other positionally licensed vowel harmony patterns. The existence of an attested first-last harmony pattern subverts the hypothesis put forth in Heinz (2018) that only SL or SP vowel harmony patterns are attested. However, if such long distance harmony patterns become local when evaluated over multi-tiered ARs, then it could be argued that the SL/SP hypothesis holds for vowel harmony. The second piece of motivation for using multi-tiered ARs is that they allow vowel harmony with transparent vowels to be SL on an autosegmental tier without relying on underspecification. Full specificataion maintains the generalization that all vowels are describable using the same set of binary features. I propose to investigate whether multi-tiered ARs predict a range of attested vowel harmony patterns. If they also do not predict unattested vowel harmony patterns I could argue that these enriched representations provide a more accurate model than strings for the cognitive computation of vowel harmony.

1.3 Autosegmental Tiers

As Rogers et al. (2013) discuss, the computational complexity of a pattern changes with different representational primitives. For example, Jardine (2019) shows that tone patterns are Autosegmental Strictly Local (ASL) when evaluated over two-tiered autosegmental representations (ARs). While some formal language theoretic work suggests that vowel harmony patterns are not SL or SP over strings, some of the vowel harmony literature utilizes ARs rather than strings. My qualifying paper was the first to unify these two approaches and suggested that vowel harmony is in fact local over multi-tiered ARs.

What is the difference between at TSL and an autosegmental tier? The TSL concept of tier consists of a subset of the segments in a word, such as non-high vowels in the case of Lokaa. Each instance of a segment on the tier is directly associated to the identical segment in the original string and thus harmony can affect only a subset of segments in a word. However, the autosegmental tiers proposed thusfar can consist of subsegmental features or tones. For vowel harmony patterns, I assume that each vowel in a word is associated to an element on multiple feature tiers because a vowel is characterized by multiple different features, i.e. height, backness, and rounding. In addition, each feature on a given autosegmental feature tier can be associated to mutiple vowels, as shown in the Lokaa example in (8).

(8) ATR harmony in Lokaa

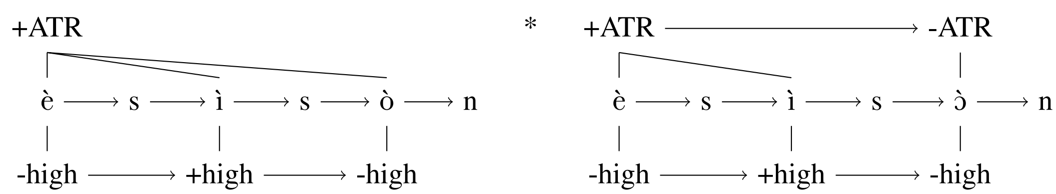

With ARs for tone, elements on a tonal tier can be multiply associated to elements on the TBU tier and vice versa. So unlike TSL tiers, autosegmental tiers consist of elements in a disjoint set from the set of elements to which they are associated and can be multiply associated. Whether it be tones multiply associated to TBUs, TBUs multiply associated to tones, or vowel features multiply associated to vowels the possibility for multiple association distinguishes autosegmental tiers from TSL tiers.

By using a more complex representation, i.e. one with tiers, the complexity of a vowel harmony pattern is reduced. I will investigate the locality of a range of vowel harmony patterns by determinging whether they can be captured using FSCs over ARs with multiple tiers. Some open questions that remain are: Where does vowel harmony fit within the complexity heirarchy of (1)? From a formal perspective, what is the range of patterns that can be represented using more than two autosegmenal tiers?

2 Autosegmental Representations

This dissertation will determine the locality of surface restrictions on vowel harmony patterns over multi-tiered ARs by investigating whether they can be captured using forbidden substructure constraints (FSCs) (Jardine, 2016, 2017). This section outlines the motivations for adopting the representations used throughout this dissertation and the basic assumptions and definitions needed to use them for analysis. I adopt a “bottlebrush” feature representation in which binary features are represented using multiple featural tiers; each feature—with a + or - value—occupies a separate tier and is associated to a vowel on the segmental tier (following Clements, 1976; Hayes, 1990; McCarthy, 1988; Padgett, 1995). Such ARs include at least one additional tier compared with the ARs of tone patterns, which utilize only two tiers (Jardine, 2016, 2017, 2019). Patterns represented with multi-tiered ARs demonstrate whether or not enriching the representation necessarily increases the expressivity of a grammar. This dissertation will determine whether or not multi-tiered ARs adequately capture vowel harmony patterns so that their expressivity can be compared to two-tiered ARs of tone.

2.1 Multi-tiered ARs

Autosegmental representations (ARs) of tonal patterns generally consist of two tiers: the TBU and segmental tiers (Goldsmith, 1976; Jardine, 2016, 2017), but an open question that remains is: From a formal perspective, what is the range of patterns that can be represented using more than two autosegmental tiers? This dissertation will investigate the expressive power needed to represent one such set of patterns. Vowel harmony patterns refer to subsegmental features, which will be represented using multiple featural tiers; each feature occupies a separate tier that is associated to a vowel on the segmental tier (following Clements, 1976; Hayes, 1990; McCarthy, 1988). For example, assuming binary features, vowel features like [± back], [± high], etc. are represented on separate tiers and associated to a vowel on the segmental tier, as in (9). Association relations are represented by straight lines that connect elements (segments and features) on different tiers. Where a tier consists of multiple elements, the successor ordering relation between elements on that tier is represented by arrows.

(9)

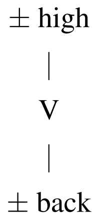

One goal for this dissertation is to extend the work of Jardine (2017) to determine whether vowel harmony patterns are local over ARs with more than two tiers, as in (9). This dissertation will also evaluate whether or not the restrictions on attested vowel harmony patterns can be captured using FSCs that contain elements of more than one feature tier.

2.2 Representational assumptions

Use of ARs requires discussion of at least some of the basic representational assumptions held throughout this dissertation. The basic assumptions are taken from Clements (1976)'s Well-Formedness Condition, which includes stipulations of *Full Specification* (FS), the *No Crossing Constraint* (NCC) (Goldsmith, 1976; Sagey, 1986), and the *Obligatory Contour Principle* (OCP) (Leben, 1973). Examples of structures that violate each of these assumptions are shown in (10)-(12) below.

(10) Violates FS

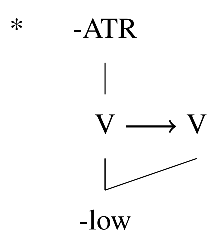

(11) Violates NCC

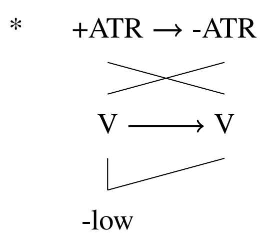

(12) Violates OCP

```
*   -ATR → -ATR
       |        |
       V ⟶ V
       |        |
    -low ⟶ -low
```

First, FS means that each featural element must be associated to at least one vowel on the segmental tier and each vowel on the segmental tier must be associated to at least one element on each featural tier. FS crucially allows vowels to be associated to multiple featural tiers as is necessary for each vowel feature to occupy its own tier. The hypothetical representation in (10) straighforwardly violates FS because there is a vowel that is not associated to any feature on the ATR tier. While both vowels are associated to a single -low feature, the second vowel is not associated to any feature on the ATR tier. Since vowel harmony patterns will be analyzed, it will be assumed that consonants cannot be associated to vowel features and that FS and vowel harmony in general ignore consonantal elements on the segmental tier.

Second, the NCC states that association lines between the segmental tier and a feature tier never cross. Odden (1994) adds that the NCC can only evaluate the association between the segmental and one featural tier at a time. The representation in (11) violates the NCC because +ATR precedes -ATR, but is associated to a vowel that is preceded by a vowel associated to -ATR; this configuration creates visually crossed association lines.

A notable effect of FS along with the NCC is that they prevent what have been called gapped structures (Archangeli & Pulleyblank, 1994; Ringen & Vago, 1998). A gapped structure is one in which a feature appears to have skipped over a vowel that it could potentially be associated to. FS would prevent gapped structures in which the "skipped" vowel is not associated to anything on that particular feature's tier. The NCC would prevent gapped structures in which the surrounding two vowels are associated to a single feature and the intervening "skipped" vowel is associated to a different feature on the same tier.

Lastly, the OCP stipulates that successive featural elements must be distinct. The representation in (12) violates the OCP because on both the ATR and low feature tiers there are two identical successive features, -ATR and -low respectively. The OCP in conjunction with FS results in representations where multiple vowels are associated to a single feature rather than having multiple successive iterations of the same feature each associated to a single vowel. An example representation of an Akan word that satisfies all of the AR properties discussed here is shown in (14).

Both the NCC and the OCP have also been derived via a concatenation operation (∘) that merges autosegmental "graph primitives"(Jardine & Heinz, 2015, p. 1). An autosegmental graph primitive consists of an element on the segmental tier, the elements on each feature tier and the associations between the featural and segmental tiers. The concatenation operation combines a finite set of adjacent graph primitives to generate a fully specified AR. For example, the AR in (14) is derived from the set of graph primitives in (13). Each primitive in (13) is concatenated with a single adjacent primitive. If two adjacent primitives share an identical feature those two features are merged into one feature with two associations, as in (14). The merging of identical adjacent features essentially prevents surface ARs from having multiple iterations of a feature and crossed associations, thus satisfying both the OCP and the NCC. However, if two segmental elements are associated to the exact same feature and a different element intervenes then both iterations of that feature will occur in the surface AR because only adjacent primitive elements are concatenated and can thus be merged. This dissertation will show that an intervening element can be a vowel associated to the same feature with a different value or a domain boundary. It will further show that a domain boundary primitive may include that boundary on both segmental and feature tiers.

(13) Concatenation of adjacent autosegmental graph primitives

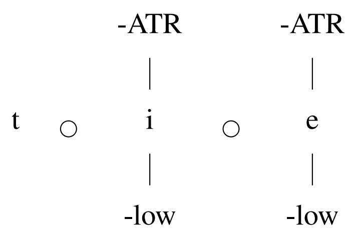

(14) Satisfies FS, NCC, and OCP

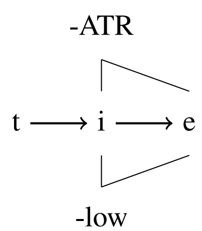

Again, the initial consonant in (14) cannot be associated to a vowel feature. While it is ordered with respect to the vowels, FS does not require the consonant to be associated to any element on either feature tier. The AR of *tie* satisfies FS because each vowel is associated to a feature on each of the featural tiers and all features are associated to at least one vowel. The AR of *tie* also satisfies both the NCC and the OCP because there is only one of each feature. The features are represented on separate tiers so association lines cannot cross and there is nothing else on those tiers that could violate the OCP. In addition, (10)-(14) illustrate that, unlike the usual notation, this paper will be adding a representation of the successor ordering relation on each tier using arrows.

2.3 Definition of Constraints

As mentioned above, this dissertation will use forbidden substructure constraints (FSCs) to determine the locality of surface restrictions on vowel harmony patterns over multi-tiered ARs. Previous work on the logical descriptions of formal languages and their applications to phonological well-formedness constraints (Heinz et al., 2011; Rogers et al., 2013) led to the development of the theory of a forbidden substructure grammar (following Jardine, 2017). A forbidden substructure grammar is a CNL of the form in (15) below. Such a grammar will generate a set of well-formed structures that does not contain any of r_1 through r_n.

(15) Forbidden substructure grammar (Jardine, 2017)
$\neg r_1 \wedge \neg r_2 \wedge \neg r_3 \wedge ... \wedge \neg r_n$

FSCs serve as a type of phonotactic restriction such that "well-formedness is based on contiguous structures of a specific size" (Jardine, 2017, p. 3). One can use FSCs as a definition of locality because they refer to elements within a structure that are connected by either an ordering or association relation. A phonological pattern is thus local if it can be described by a forbidden substructure grammar because it can be captured with FSCs by referring to a subset of the elements within structures and their connections. Jardine (2017) uses FSCs to show that attested tone patterns are local in this way. This dissertation will utilize FSCs over multi-tiered ARs to show that a variety of vowel harmony patterns are local in the same way.

FSCs are written using an asterisk ($*$) to denote that a piece of structure is forbidden, or marked. Over multi-tiered ARs, the forbidden piece of structure includes elements on one or multiple tiers and their connections, as in (16). Vowel harmony patterns can utilize a feature on one tier in order to constrain the harmony that occurs on another tier. For example, the FSC in (16) would restrict harmony in the feature F such that no structure can contain more than one feature on the F feature tier without the presence of a -G feature.

(16) Example of a FSC over multi-tiered ARs

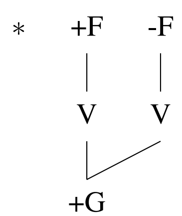

The FSC in (16) captures a pattern with harmony in the feature F and blocking of F-harmony by a -G feature. Example structures of hypothetical words, excluding consonants, are shown in (17).

(17) F harmony with blocking -G vowel

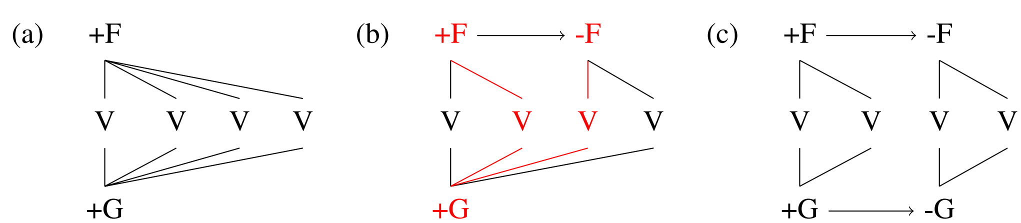

The hypothetical word in (17a) illustrates full harmony in both features F and G and it is grammatical because it does not contain the forbidden substructure. The word in (17b) has features with two different values on the F tier, but only one G feature; it contains the forbidden substructure in red and is thus ungrammatical. Lastly, the word in (17c) would be grammatical because it contains two different F features and a -G feature, so it does not contain the forbidden substructure. Concrete examples of FSCs and how they are used to capture harmony patterns will be shown in section 4.

3 Theory Proposal: Vowel harmony without correspondence

Vowel harmony is an assimilation process that affects only the vowels in a word. Previous work has analyzed vowel harmony patterns as resulting from a single assimilation process, whether it be feature spreading or agreement (Bakovic, 2000; Clements, 1976; McCarthy, 2011; Nevins, 2010; Rose & Walker, 2011; van der Hulst & Smith, 1986; Walker, 2010). To begin, I will investigate vowel harmony as a phonotactic restriction rather than a process that changes an input form into an output. As a theory of phonotactics, FSCs restrict only the possible output surface ARs in a language. For example, spreading is traditionally understood as a process that changes an input into an output with multiple vowels associated to a single feature. Thus, on the surface, a spreading AR is one with multiple association, as shown in (18a). Similarly, agreement is traditionally understood as a process that changes an input into an output surface form with multiple vowels associated to different iterations of a feature that have the same value. A surface agreement AR is one in which vowels are associated to different iterations of a feature with the same value, as in (18b). In short, "agreement" means that vowels don't have differing values for a particular feature. Following Walker (2011), this dissertation will show that given a uniform theory of surface markedness constraints vowel harmony patterns utilize ARs, as defined in the previous section, that can reflect either type of assimilation.

(18) Surface ARs of assimilation mechanisms

a. Spreading b. Agreement

Walker (2011)'s theory of feature licensing accounts for a variety vowel harmony patterns with different assimilation mechanisms. She utilizes ARs similar to the ones defined in the previous section and her licensing patterns parallel the assimilation mechanisms in (18). For example, indirect licensing is a pattern wherein harmony is the result of licensing, which spreads a feature, and so the resulting surface AR utilizes multiple association, as in (18a). Similarly, identity licensing is a pattern in which harmony is the result of licensing between two non-adjacent and corresponding features, as in (19). Walker (2011) further observes that languages with vowel harmony via identity licensing also necessarily have harmony via indirect licensing (i.e.in words with no transparent vowels), but the reverse implication does not hold.

(19) Surface AR of Walker (2011)'s identity licensing

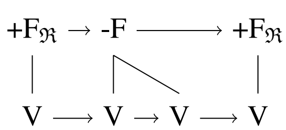

The difference between the agreement AR in (18b) and Walker (2011)'s identity licensing configuration in (19) is the addition of a correspondence relation between harmonic features in the latter, represented with a subscript $\mathfrak{R}$. In identity licensing patterns, non-adjacent harmonizing features necessarily corresponds because they are separate instances of a single feature.

I propose that vowel harmony patterns, which utilize agreement are SL over multi-tiered ARs because agreeing features are connected to the same intervening feature by the successor ordering relation regardless of how many segments are associated to that intervening feature. Thus on the surface, harmony propagation is demonstrated between non-adjacent features on a tier because they have the same + or - value, which differs from the intervening feature value on that same tier; and there is no need to connect identical features with an additional correspondence relation. In addition, I corroborate Walker (2011)'s observation that languages with vowel harmony via agreement also utilize feature spreading, such as in words with no transparent vowels. Thus, vowel harmony can be the result of more than one assimilation process, even within the same language.

4 Case Studies

This section demonstrates the locality of vowel harmony with neutral vowels via each of the two assimilation mechanisms in (18). ATR harmony with blocking vowels in Akan exemplifies vowel harmony via feature spreading and can be captured by a single FSC over a multi-tiered AR. Back harmony with transparent vowels in Finnish exemplifies vowel harmony via agreement and can be captured by a set of FSCs over multi-tiered ARs.

This dissertation constitutes the second formal language theoretic study of vowel harmony as a phonotactic restriction, following my qualifying paper. It will take a slightly different approach than has been taken before by evaluating only the restrictions on output substructures. While vowel harmony has been considered a derivational process, the aim here is to determine the locality of only the surface restrictions on vowel harmony patterns over multi-tiered ARs. The harmonizing ARs that will be examined contain at least one feature that is associated to more than one vowel, as it would be on the surface. Ignoring input structures in this way allows for the eventual classification of vowel harmony within the sub-SF hierarchy of patterns, which in turn allows for the comparison of vowel harmony with other phonological patterns that have been classified on the same hierarchy, such as tone in Jardine (2019).

The case studies presented here are analyses of vowel harmony as a phonotactic restriction rather than a process that changes an input form into an output. Phonotactics restrict the possible forms that can be found on the surface regardless of the input, so this section presents analyses of only surface vowel harmony patterns. Because these case studies take a different perspective than some traditional accounts of vowel harmony, they utilizes a slightly different understanding of certain terminology as well. As mentioned above, vowel harmony is traditionally considered an assimilatory process, but in the next section only surface forms will be analyzed. Spreading only refers to the output surface form and a spreading AR is one in which multiple vowels are associated to a single feature (i.e."spreading" means multiple association). In addition, agreement refers only to the output surface form so an agreement AR is one in which vowels are associated to different iterations of a feature with the same value. In short, "agreement" means that vowels don't have differing values for a particular feature. This section thus adapts the terminology traditionally used to discuss vowel harmony as a transformational process and uses the terms to refer only to output structures that can be restricted by the phonotactic constraints of a language.

In this section, I show that two vowel harmony patterns with neutral vowels are local over multi-tiered ARs because they can be captured with FSCs. In 4.1, I present ATR spreading with a blocking low vowel in Akan. In 4.2, I present back agreement with transparent vowels in Finnish.

4.1 Blocking vowels

An example of vowels that block harmony is found with ATR harmony in Akan (Clements, 1976). The Akan vowel inventory, in Table 1, consists of ten vowels with two main featural distinctions: ± ATR and ± low. There are two +low vowels, [ɜ] and [a], +ATR and -ATR, respectively. All other vowels are considered -low and distinguished by ATR such that the +ATR vowels are [i, e, u, o] and the -ATR vowels are [ɪ, ɛ, ʊ, ɔ].

Table 1
Akan Vowels

	+ATR	-ATR
-low	i	ɪ
	u	ʊ
	e	ɛ
	o	ɔ
+low	ɜ	a

The harmony generalization is that if a word contains a sequence of -low vowels, then those vowels will also share the same ATR feature (Clements, 1976). For example, the words in (20) contain only -low vowels, which are also all either +ATR or -ATR.

(20) -low vowels share an ATR feature value
  a. tie ‘listen’
  b. obejii ‘he came and removed it’
  c. ɔbɛjɛɪ ‘he came and did it’
  d. wubenumˀ‘you will suck it’
  e. wʊbɛnʊmˀ‘you will drink it’

The surface requirement that -low vowels share the same ATR feature can also be written as a FSC, which forbids two vowels associated to the same -low feature from being associated to different ATR features, as in (21). The ordering relation on the ATR tier in (21) is omitted because the + or - values of the two ATR features are irrelevant for this constraint, as long as they differ. The ordering relation on the segmental tier of this FSC is also omitted and the reason will be made clear by the example in (22).

(21)

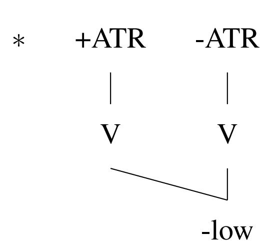

(22) [obejii] 'he came and removed it'

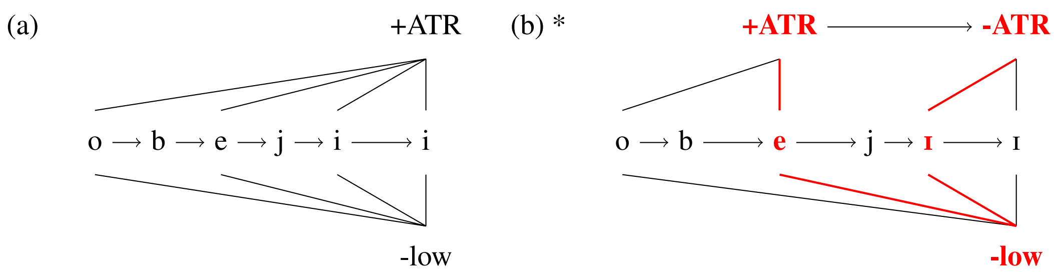

The AR for the grammatical Akan word [obeijii] "he came and removed it" is shown in (22a). Here a single +ATR and a single -low feature are each associated to each vowel within the word, demonstrating full ATR and low harmony. On the other hand, the hypothetical Akan word, [obejɪɪ], represented in (22b) is ungrammatical because it demonstrates full -low harmony, but does not demonstrate full ATR harmony; so, the AR in (22b) contains the forbidden structure of (21), shown in bold and red.

However, in traditional vowel harmony terms the presence of a +low vowel blocks the rightward spread of ATR, some examples are shown in (23). Translating this to the static surface representations assumed here, two -low vowels must be associated to the same ATR feature, but if a +low vowel intervenes they can be associated to different ATR features. The representation of (23a) exemplifies this pattern and is shown in (24).

(23) Vowels on either side of +low can have different ATR features

a. pɪrɜko 'pig'
b. obisaɪ 'he asked'
c. mɪkɔkɜri 'I go and weight it'
d. okog^{w}arɪʔ 'he goes and washes'

(24) [pɪrɜko] 'pig'

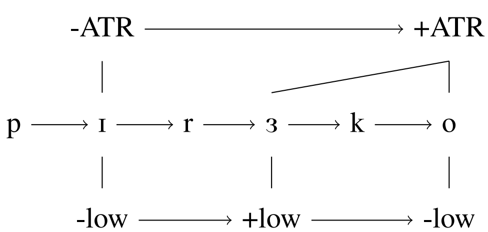

Crucially, the AR in (24) does not contain the FSC from (21). While the AR for [pɪrɜko] "pig" does contain two vowels associated to a -low feature and two different ATR features, they are each separately concatenated to the intervening [ɜ] vowel, which is associated to a +low feature. So, the surrounding vowels are associated to two separate -low features on the surface and thus the AR satisfies FS and the NCC. Because the forbidden structure is not present [pɪrɜko] "pig" is grammatical.

In summary, the vowel harmony pattern with blocking vowels in Akan can be captured using the FSC in (21), which does not refer to the successor relation on any tier. Akan vowel harmony could thus be considered local because the FSC that captures the pattern need only refer to the associations between vowels and features. The next section outlines a vowel harmony pattern with transparent vowels.

4.1.1 Surface spreading is local. Some previous analyses of vowel harmony assume that all harmony patterns result from a single assimilation process: feature spreading. Feature spreading is generally considered to be a transformation from an underlying representation in which a single feature is associated to a single vowel into a surface representation with multiple vowels associated to the same feature, as in (25). In other words, feature spreading maps an underspecified underlying AR onto a fully specified surface AR with multiple association.

(25) Surface spreading

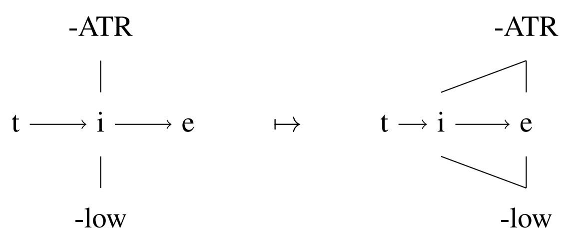

This section focuses only on surface representations and Akan provides an example of a pattern in which vowel harmony assimilation is represented by spreading ARs. The surface spreading ARs used throughout this paper consist of a single feature that is associated to multiple vowels. Akan provides an example of a classic spreading pattern, in which an initial vowel feature (ATR) is associated to all the vowels in a word to the left of a +low blocking vowel, as shown in (22a).

The analysis of Akan provided here demonstrates that spreading ARs are local on the surface. Here locality means that spreading ARs consist of a domain defined by a single ATR feature node, they must include a contiguous span of vowels, but they are not bounded in length, as in (22a); or when two different ATR features are present, one succeeds the other regardless of how many vowels are associated to each. In addition, the FSC posited for Akan is able to capture the Akan ATR harmony pattern for words with and without blocking vowels.

4.2 Transparent vowels

Finnish provides an example of backness harmony with four transparent vowels. The Finnish vowel inventory in Table 2 consists of 16 vowels with contrastive length and three main featural distinctions: ± back, ± low, and ± round (Ringen & Heinamaki, 1999; Välimaa-Blum, 1986). The four vowels transparent to backness harmony, [i, iː, e, eː], are all [-back, -round, -low]. Of the harmonizing vowels [y, yː, u, uː, ø, øː, o, oː] are all +round and -low while [æ, æː, ɑ, ɑː] are all +low and -round. The +back vowels are [u, uː, o, oː, ɑ, ɑː] and the -back vowels are [i, iː, e, eː, y, yː, ø, øː, æ, æː]. The difference between harmonizing and transparent Finnish vowels is characterized by low and round feature values. Transparent vowels are all [-low, -round] and thus harmonizing vowels have a positive value for the low and/or round feature.

Table 2

Finnish Vowels

	-round	+round		
-low	i, iː	y, yː	u, uː	
	e, eː	ø, øː	o, oː	
+low		æ, æː	ɑ, ɑː	-round
	-back		+back	

The Finnish harmony generalization is that all of the harmonizing vowels in a root will share the same back feature with each other and harmonizing suffix vowels will share the same back feature with the harmonizing root-final vowel (Nevins, 2010; Ringen & Heinamaki, 1999; van der Hulst, 2017; Välimaa-Blum, 1986). Since the same harmony generalization holds for both root and suffix vowels the Finnish generalization can also be stated as two harmonizing vowels must share the same back feature. For example, the words in (26) contain only +round or +low vowels, which are also either all +back or all -back.

(26) harmonizing vowels share a back feature value

a. pøytæ 'table'
b. kæntæː 'turn'
c. tykætæ 'like'
d. poutɑ 'fine weather'
e. murtɑː 'break'
f. kokɑtɑ 'cook'

Transparent vowels, however, do not block or undergo harmony so in the Finnish words in (27) +back harmony appears to skip over the [-back, -round, -low] vowels [i, iː, e, eː]. The novel contribution of the current analysis is to treat transparent vowels in the same way as harmonizing vowels; the FSCs posited in this section are able to generate the Finnish pattern without underspecification of back features.

(27) back harmony skips over transparent vowels

a. ruvetɑ 'start'
b. tuoliɑ 'chair'
c. lukeɑ 'read (inf.)'
d. kɑuneus 'beauty'
e. nɑivius 'naiveness'
f. kotikɑs 'cozy'

The surface requirement that +round and +low vowels share the same back feature can also be stated negatively as a constraint that forbids either a +round or a +low vowel from being associated to a different preceeding back feature. Together, the four FSCs in (28) generate this negative constraint and the Finnish vowel harmony pattern. The ordering relation on the segmental tier of the FSCs is omitted because the vowels can have consonants between them, as in (29). The ordering relation on the back tier, however, is crucial in order to allow transparency of certain -back vowels; all elements in an FSC must be connected.

(28)

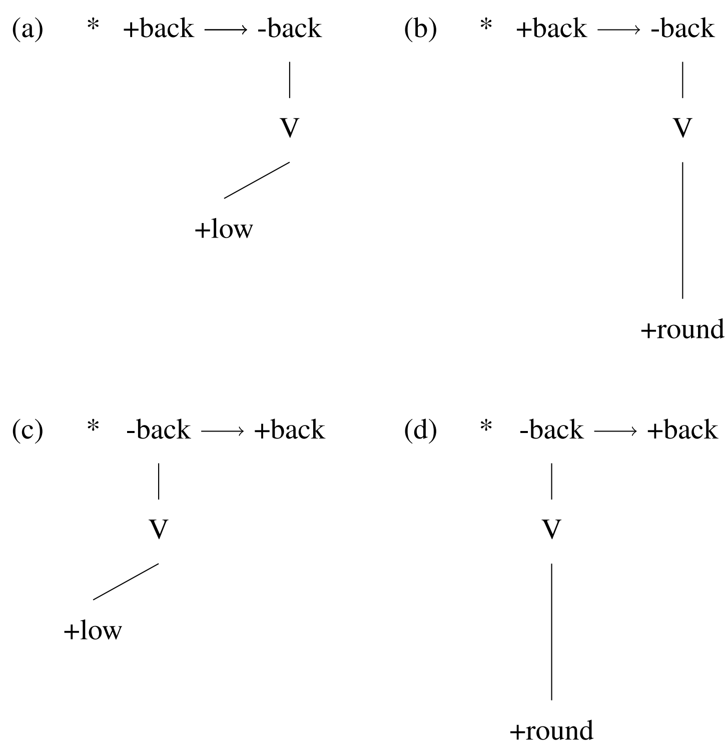

The ARs in (29) illustrate how the Finnish FSCs rule out ungrammatical disharmonic words. The AR for the grammatical Finnish word [poutɑ] ‘fine weather’, shown in (29a), contains both a +round and a +low non-initial vowel as well as a single +back feature, which demonstrates full back harmony. The hypothetical Finnish word, [poutæ] in (29b), however, contains the forbidden structure of (28a) in bold and red. In (29b) the final vowel does not harmonize with the penultimate vowel because they are associated to different back features.

(29) [poutɑ] ‘fine weather’

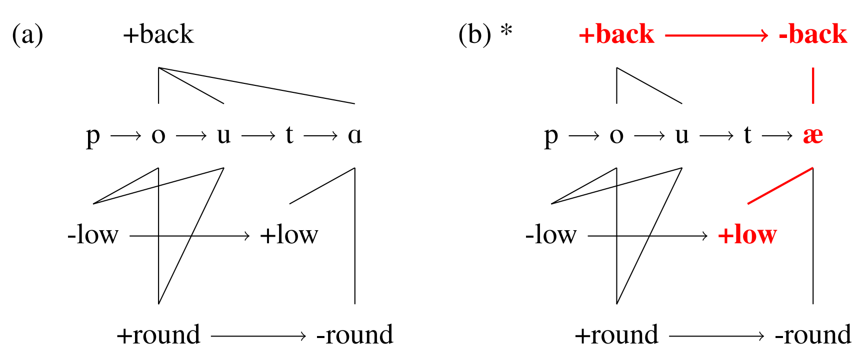

Crucially, the behavior of transparent vowels with respect to vowel harmony in Finnish is captured by the four FSCs in (28) without reference to underspecification of back features. For example, the words in (27) all contain vowels with -back features that follow +back vowels, but because the -back vowels are also [-low, -round] the words are grammatical. The transparent vowels are associated to features on the same tiers as the harmonizing vowel features and their transparency results from the interaction of the -back features with -low and -round features, as shown in (30). Because the Finnish FSCs only forbid associations to -back features when vowels are also either +low or +round, the [-back, -low, -round] vowels are able to occur anywhere within a word. While Finnish does have [-back, +low] and [-back, +round] vowels, they do not occur unless all the vowels in a word are associated to a single -back feature because the Finnish FSCs only forbid [-back, +low] and [-back,+round] vowels when the -back feature is either preceeded or succeeded by a +back feature. This additional restricion enforces +back agreement across transparent vowels; it is only ever the case that a +back and a -back feature are in a successor relation if the -back vowel is also -low and -round. So, in words with more than one back feature any -back vowel must be transparent and all other vowels must be +back.

(30) [ruvetɑ] 'start'

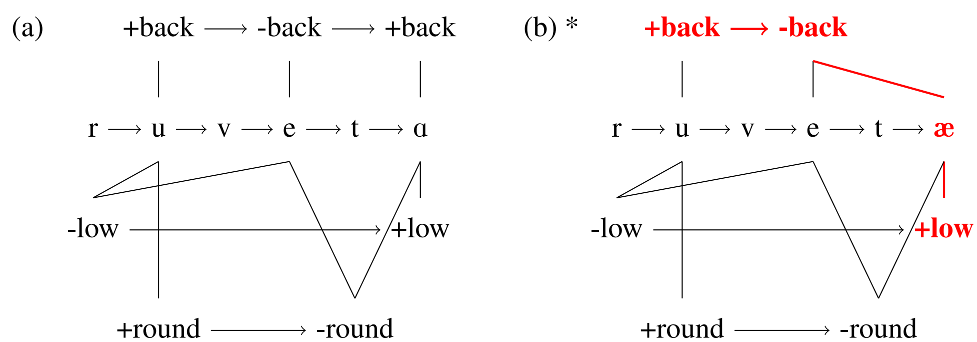

In (30a) the [u] and [ɑ] vowels are each associated to a +back feature, but not the same one. The [e] vowel occurs between them and is associated to a -back feature. The two +back features are also in a successor relation with the intervening -back feature. The AR in (30a) is grammatical because the -back vowel is not associated to a +low or a +round feature, so the AR does not violate any of the FSCs in (28). The AR in (30b), on the other hand, contains a [-back, +low] vowel, and so (28a) is violated, as shown in bold and red. Despite being separated by a transparent vowel, it is still necessary for the suffix and root vowels to agree and the same FSCs that capture Finnish back harmony in (30a) also enforce agreement across a transparent vowel by marking words like (30b) as ungrammatical.

In summary, the vowel harmony pattern with transparent vowels in Finnish can be captured using the four FSCs in (28). These FSCs refer to the successor relation(s) on the back tier, which also interacts with both the round and low feature tiers. Finnish vowel harmony could thus be considered local because the FSCs that capture the pattern refer to the associations between vowels and features and the ordering between features.

4.2.1 Surface agreement is local. Finnish exemplifies an assimilation mechanism that differs from spreading. In Finnish there are grammatical words like (30a), which contain two +back harmonizing vowels with a transparent vowel between them. The NCC prevents Finnish from using spreading ARs because a single +back feature cannot be associated to a vowel across an intervening -back feature. Two +back features can occur because they are not in a successor relation with each other, so ARs do not violate the OCP. So, the assimilation between two vowels of a +back feature in words like (30a) must be due to a mechanism that differs from feature spreading because the two +back features are not successive. In addition, the OCP allows multiple iterations of a +back feature to occur as long as each is in a successor relation with the intervening -back feature. In this paper, this other type of assimilation is called *agreement*. Agreement is represented on the surface as an AR in which two non-successive features on a tier share a value and the intervening feature on that tier has the opposite value, as shown in (30) and (31).

(31) [mɑisemiɑ] 'scenery.plural.partitive'

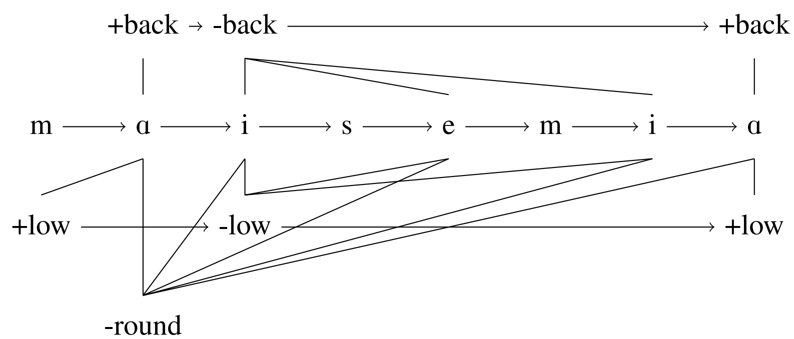

The analysis of Finnish provided here demonstrates that agreement ARs are local on the surface. Here locality means that agreement ARs consist of a domain defined by a single -back feature, which is associated to a contiguous span of vowels of unbounded length, and which both precedes and succeeds +back features. Identical +back features are connected via the successor relation to the single -back feature between them regardless of the number of vowels associated to the intervening -back feature. For example, in (31) more than one transparent vowel is associated to a -back feature that intervenes between two +back features. On the segmental tier it would appear that two +back vowels, such as [ɑ] and [o], can be separated by more than one -back vowel, but on the back feature tier the +back and -back features are in a successor relation. The +back agreement is local because it appears to skip over un unbounded number of transparent vowels, which are all associated to a single -back feature that is also in a successor relation with the agreeing +back features. The grammatical AR in (31) also does not violate the Finnish FSCs in (28) and so the Finnish FSCs are still able to capture the agreement pattern. The successor relation on the back tier further allows transparent vowels to be associated to a feature on the same tier as harmonizing vowels, rather than being underspecified. Thus Finnish vowel harmony demonstrates surface agreement and can be considered local.

5 Proposal: Further Investigations

The theory outlined in the previous section analyzes vowel harmony as a phonotactic restriction using FSCs over multi-tiered surface ARs. This theory has also been extended to account for two morphologically conditioned harmony patterns in Turkish by adding a morpheme boundary primitive to the set of possible representations, which includes the boundary on all feature tiers as well as on the segmental tier. In addition a hypothetical logically possible sour grapes pattern was captured with FSCs over multi-tiered ARs enriched with a word boundary. I propose to investigate the possible extensions of the theory presented above in order to determine the locality of a wide variety of attested vowel harmony patterns as well as the logically possible patterns it predicts.

In order to compare the theory presented here with Walker (2011)'s licensing theory, I will need to verify that my theory can capture the same patterns. I will start by reanalyzing the simplest vowel harmony patterns accounted for by Walker (2011)'s identity licensing. In this way, I will determine whether the two theories differ in their basic predictions for transparency. Walker (2011) analyzes a wide variety of languages and two patterns of interest that I will highlight here are Eastern Meadow Mari and Baiyina Orochen. The vowel harmony pattern in Eastern Meadow Mari is described as a first-last harmony pattern. Eastern Meadow Mari vowels contrast in backness, height, and rounding. There are three different suffixes, in which the vowel alternates in backness depending upon the back value of the root-initial vowel: nominative singular first person plural possessive [næ/na], nominative singular second person plural possessive [tæ, dæ/ta], and the dative [læn/lan].

Table 3
Eastern Meadow Mari Vowels

	-back		+back	
+high	i	y		u
-high	e	ø	ə	o
	æ		a	
	-round	+round	-round	+round

(32) Nom.sg.1stpl poss. suffix (Vaysman, 2009; Walker, 2011)
  a. ij-næ 'our year'
  b. ʃot-na 'our sense'
  c. køgørtʃen-næ 'our dove'
  d. pusaŋgə-na 'our tree'

(33) Nom.sg.2ndpl poss. suffix (Vaysman, 2009; Walker, 2011)
  a. em-dæ 'your(pl) medecine'
  b. tʃødræ-tæ 'your(pl) forest'
  c. kutko-ta 'your(pl) ant'

(34) Dative suffix (Vaysman, 2009; Walker, 2011)
  a. pel-læn 'half (dat.)'
  b. sør-læn 'milk (dat.)'
  c. lum-lan 'snow (dat.)'
  d. keneʒ-læn 'summer (dat.)'
  e. korno-lan 'road (dat.)'

The data in (32-34) exhibit full harmony, but there are three particular vowels in Eastern Mari that behave almost like transparent vowels: [ə], [a], and [e]. In (35a-b) below a [+back] [ə] occurs between two [-back] vowels and in (35c) a [+back] [a] occurs between two [-back] vowels.

(35) a. yremə-næ 'our street'
b. tyrə-tæ 'your(pl) edge'
c. meraŋ-læn 'hare (dat.)'

However, in (36) when these two [+back] vowels are in the initial syllable they occur with a [+back] suffix vowel. So [ə] and [a] appear to participate in harmony when in an initial syllable, but behave as transparent when non-initial.

(36) a. tam-na 'our taste'
b. pərəs-na 'our cat'
c. aβam-ta 'your (pl) mother'
d. rəwəʒ-ta 'your (pl) fox'
e. pasa-lan 'work (dat.)'
f. təlzə-lan 'moon (dat.)'

It could be argued that the [+back] suffix vowel is the default and so it occurs with the transparent [+back, -high, -round] vowels, but there is also a [-back] vowel with the same behavior: [e].

(37) a. em-dæ 'your (pl) medecine'
b. keneʒ-læn 'summer (dat.)'
c. pel-læn 'half (dat.)'
d. meraŋ-læn 'hare (dat.)'

(38) a. pareŋə-na 'our potato'
b. uβer-ta 'your (pl) news'
c. təlze-ta 'your (pl) moon, month'

In (37) the suffix vowel surfaces as [-back] when the root vowel(s) are [e]. Even across a [+back] vowel, the suffix and initial vowel have the same back specification. However, The words in (38) demonstrate that the [e] vowel can also behave as transparent when it occurs between two [+back] vowels. While the [-back, -high, -round] [e] sometimes behaves as a transparent vowel, the same is not true of the [-back, -high, -round] [æ] vowel. Lastly, the three vowels with transparent-like behavior ([ə], [a], and [e]) do not form a natural class. Walker (2011) thus describes the Eastern Meadow Mari generalization as one in which a suffix vowel harmonizes only with the initial vowel.

Walker (2011)'s description has consequences for computational theory, though. The Eastern Meadow Mari pattern outlined above is described as first-last harmony. Heinz (2018) shows that first-last harmony is LT over strings and so a theory of phonotactics as only SL, SP, or TSL predicts this pattern to be unattested. However, if Eastern Meadow Mari's harmony pattern is local over multi-tiered ARs, it could be argued that the theory proposed here more accurately captures the range of attested vowel harmony patterns than these string-based classes. I will determine the complexity of Eastern Meadow Mari vowel harmony by finding out if it is possible to capture this pattern with FSCs over multi-tiered ARs.

Another interesting pattern is found in Baiyina Orochen, which is described as having rounding harmony triggered only by root-initial non-high short vowels and blocked by high vowels. Vowel length is not generally considered to be a feature of vowel quality and so it is not usually considered to participate in vowel harmony. In order to account for Baiyina Orochen rounding harmony I will have to determine whether vowel length should be considered as a vowel quality feature on its own tier in the proposed theory and what predictions that might make. Additional languages that Walker (2011) captures with licensing and which I will analyze with FSCs over multi-tiered ARs include: Buchan Scots, Macuxi, Ticinese, C'Lela, Old High German, Jaqaru, Chamorro, Vinalopo Mitja, Servigliano, Hungarian, Francavilla Fontana, Lena, Ascrea, and Eastern Andalusian.

A goal of this dissertation is to situate vowel harmony within the complexity hierarchy in (1). Over string representations, attested vowel harmony patterns have been claimed to be TSL (Aksënova, 2017). More complex harmony patterns, such as first-last harmony, have been claimed to be unattested (Heinz, 2018), but Walker (2011) captures a first-last harmony pattern in Eastern Meadow Mari. Following Jardine(2016, 2017, 2019), I propose to determine whether or not there is an existing formal language complexity class that encompasses the range of attested vowel harmony patterns over multi-tiered ARs.

Lastly, many scholars have viewed vowel harmony as mapping an input with a vowel feature associated to one vowel onto an output where that same feature is associated to multiple vowels. However, a hierarchy that classifies sets of ARs—based on the Chomsky and related subregular hierarchies—differs significantly from a parallel hierarchy for sets of pairs of ARs, such as in a transformation (or map) from underlying to surface form. Some influential work has been dedicated to classifying input-output maps in phonology as stricly local within a Chomsky-based hierarchy of sets of pairs of strings and demonstrating their learnability (Chandlee & Heinz, 2018; Chandlee & Jardine, 2013, 2019; Chandlee, Eyraud, & Heinz, 2014). Other recent work uses automata theory to argue that certain pathological harmony and tone spreading patterns are more complex because they require weakly deterministic transductions (Heinz & Lai, 2013; McCollum, Bakovic, Mai, & Meinhardt, 2018). In this dissertation I will utilize the newly developed quantifier-free least-fixed point logic to investigate how local surface phonotactic constraints over multi-tiered ARs inform a transformational analysis of vowel harmony patterns (Chandlee & Jardine, 2019; Koser et al., 2019).

6 Tentative Outline of Dissertation and Timeline for Completion

6.1 Outline

1. Introduction
    - 1.1 Formal Language Theory and Phonology
    - 1.2 Locality
2. Surface Vowel Harmony Patterns
    - 2.1 Autosegmental Representations
    - 2.2 Previous Theories of Vowel Harmony
        - 2.2.1 Spreading (Clements, 1976; McCarthy, 2011)
        - 2.2.2 Agreement (Bakovic, 2000)
        - 2.2.3 Licensing (Walker, 2011)
    - 2.3 Assimilation Mechanisms
        - 2.3.1 Spreading
        - 2.3.2 Agreement
    - 2.4 Terminology
3. Computational Complexity
    - 3.1 Subregular Hierarchy
    - 3.2 Logical Descriptions of Patterns
    - 3.4 Forbidden Substructure Grammars
4. Spreading Analyses
    - 4.1 Unbounded Spreading
    - 4.2 Blocking Vowels
        - 4.2.1 Akan
5. Agreement Analyses
    - 5.1 Transparent Vowels
        - 5.1.1 Finnish
    - 5.2 First-Last Harmony
        - 5.2.1 Eastern Meadow Mari
6. Bounded (Non-iterative) Harmony
7. Logical Transductions as Input-Output Maps
    - 7.1 Quantifier-Free Least Fixed Point Logic
    - 7.2 Unbounded Spreading
    - 7.3 Bounded Harmony
    - 7.4 Transparency
8. Future Investigations
9. Conclusion

6.2 Timeline

Summer 2019

May 2019-Write Phonology Seminar term paper: Extend QFLFP to describe multi-tiered ARs
June 2019-FSC Analysis: Eastern Mari, Baiyina Orochen
July 2019-FSC Analysis: Unbounded Spreading, Blocking
August 2019-FSC Analysis: Transparent vowels

Fall 2019

I will be teaching in the writing program, so I am planning not to have time for research during the fall semesters. In addition, Adam will be on sabbatical this semester so this work will be on hold. However, the writing program guarantees an extra year of funding and I will have spring semesters free to continue the dissertation.

Spring 2020

January 2020-FSC Analysis: Bounded Patterns
February 2020-Transformational Analysis: Unbounded spreading, Blocking
March 2020-Transformational Analysis: Bounded spreading
April 2020-Transformational Analysis: Transparent vowels

Summer 2020

May 2020-Transformational Analysis: Eastern Meadow Mari
June 2020-Transformational Analysis: Baiyina Orochen
July 2020-Write chapter 2, section 2.1-2.2
August 2020-Finish chapter 2

Fall 2020

Unless I receive some other funding I will be teaching in the writing program again this semester. As in the previous year, I will plan not to have time for dissertation work during this semester.

Spring 2021

January 2021-Write chapter 3
February 2021-Write chapter 8-9
March 2021- Write Introduction
April 2021- Defend

7 References

Akinlabi, A. (2009). Neutral vowels in lokaa harmony. *Canadian Journal of Linguistics*, *59*(2), 197–228.

Aksënova, A. (2017, April). Tier alphabet(s): Tendencies and restrictions. Qualifying paper.

Aksënova, A., & Deshmukh, S. (2018). Formal restrictions on multiple tiers. In *Proceedings of the society for computation in linguistics* (Vol. 1, p. Article 8).

Archangeli, D., & Pulleyblank, D. (1994). *Grounded phonology* (Vol. 25). MIT Press.

Bakovic, E. (2000, January). *Harmony, dominance, and control* (PhD thesis). Rutgers University, New Brunswick, New Jersey.

Chandlee, J., & Heinz, J. (2018). Strict locality and phonological maps. *Linguistic Inquiry*, *49*(1), 23–60.

Chandlee, J., & Jardine, A. (2013). Learning phonological mappings by learning strictly local functions. In *Proceedings of the 2013 annual meeting on phonology*.

Chandlee, J., & Jardine, A. (2019). Autosegmental input strictly local functions. In *Transactions of the association for computational linguistics* (Vol. 7, pp. 157–168).

Chandlee, J., Eyraud, R., & Heinz, J. (2014). Learning strictly local subsequential functions. In *Transactions of the association for computational linguistics* (Vol. 2, pp. 491–503).

Clements, G. (1976). Vowel harmony in non-linear generative phonology: An autosegmental model. Bloomington, Indiana University Linguistics Club.

Goldsmith, J. (1976). *Autosegmental phonology* (PhD thesis). Massachusetts Institute of Technology.

Hayes, B. (1990). Diphthongization and coindexing. *Phonology*, *7*(1), 31–71.

Heinz, J. (2018). The computational nature of phonological generalizations. In L. Hyman & F. Plank (Eds.), *Phonological typology* (Vol. 23, pp. 126–195). Mouton de Gruyter.

Heinz, J., & Idsardi, W. (2013). What complexity differences reveal about domains in language. *Topics in Cognitive Science*, *5*(1), 111–131.

Heinz, J., & Lai, R. (2013). Vowel harmony and subsequentiality. In *Proceedings of the 13th meeting on the mathematics of language (mol 13)* (pp. 52–63). Association for Computational Linguistics.

Heinz, J., Rawal, C., & Tanner, H. G. (2011). Tier-based strictly local constraints for phonology. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies: Short papers* (Vol. 2). Association for Computational Linguistics.

Jardine, A. (2016). *Locality and non-linear representations in tonal phonology* (PhD thesis). University of Delaware.

Jardine, A. (2017). The local nature of tone association patterns. *Phonology*, *34*(2), 385–405.

Jardine, A. (2019). The expressivity of autosegmental grammars. *Journal of Logic, Language, and Information*, *28*(1), 9–54.

Jardine, A., & Heinz, J. (2015). A concatenation operation to derive autosegmental graphs. In *Proceedings of the 14th annual meeting on the mathematics of language (mol 2015)* (pp. 139–151). Chicago, USA: Association for Computational Linguistics.

Koser, N., Oakden, C., & Jardine, A. (2019). Tone association and output locality in non-linear structures. In K. Hout, A. Mai, A. McCollum, S. Rose, & M. Zaslansky (Eds.), *Proceedings* *of the 2018 annual meeting on phonology*. Washington DC: Linguistics Society of America.

Lai, R. (2015). Learnable vs. unlearnable harmony patterns. *Linguistic Inquiry*, *46*(3), 425–451.

Leben, W. (1973). *Suprasegmental phonology* (PhD thesis). Massachusetts Institute of Technology.

McCarthy, J. (1988). Feature geometry and dependency: A review. *Phonetica*, *38*. Retrieved from http://scholarworks.umass.edu/linguist_faculty_pubs/38

McCarthy, J. (2011). Autosegmental spreading in optimality theory. In *Tones and features* (Clements Memorial Volume., Vol. 27). Retrieved from https://scholarworks.umass.edu/linguist_faculty_pubs/27

McCollum, A., Bakovic, E., Mai, A., & Meinhardt, E. (2018, September). *The expressivity of segmental phonology and the definition of weak determinism*.

Nevins, A. (2010). *Locality in vowel harmony*. *Linguistic Inquiry Monographs* (Vol. 55). MIT Press.

Odden, D. (1994). Adjacency parameters in phonology. *Language*, *70*(2), 289–330.

Padgett, J. (1995). Feature classes. In J. Beckman, S. Urbanczyk, & L. Walsh (Eds.), *Papers in optimality theory* (Vol. 18, pp. 385–420).

Padgett, J. (2002). Feature classes in phonology. *Language*, *78*(1), 81–110. Retrieved from http://www.jstor.org/stable/3086646

Ringen, C., & Heinamaki, O. (1999). Variation in finnish vowel harmony: An ot account. *Natural Languge and Linguistic Theory*, *17*, 303–337.

Ringen, C., & Vago, R. (1998). Hungarian vowel harmony in optimality. *Phonology*, *15*, 393–416.

Rogers, J., & Pullum, G. (2011). Aural pattern recognition experiments and the subregular hierarchy. *Journal of Logic, Language, and Information*, *20*, 329–342.

Rogers, J., Heinz, J., Fero, M., Hurst, J., Lambert, D., & Wibel, S. (2013). Cognitive and sub-regular complexity. *Formal Grammar*, 90–108.

Rose, S., & Walker, R. (2011). Harmony systems. In J. Goldsmith, J. Riggle, & A. Yu (Eds.), *The handbook of phonological theory* (pp. 240–290). Blackwell.

Sagey, E. (1986). *The representation of features and relations in non-linear phonology* (PhD thesis). Massachusetts Institute of Technology.

van der Hulst, H. (2017). A representational account of vowel harmony in terms of variable elements and licensing. In *Approaches to hungarian* (Vol. 15). John Benjamins Publishing Company.

van der Hulst, H., & Smith, N. (1986). On neutral vowels. In *The phonological representation of suprasegmentals* (pp. 233–281).

Välimaa-Blum, R. (1986). Finnish vowel harmony as a prescriptive and descriptive rule: An autosegmental account. In F. Marshall (Ed.), *Proceedings of the third eastern states conference on linguistics*. University of Pittsburgh.

Walker, R. (2010). Nonmyopic harmony and the nature of derivations. *Linguistic Inquiry*, *41*(1), 169–179.

Walker, R. (2011). *Vowel patterns in language* (Vol. 130). Cambridge University Press.

Walker, R. (2014). Surface correspondence and discrete harmony triggers. In *Proceedings of the annual meeting on phonology*.