

THE EFFECTS OF NON-LINEAR DATA STRUCTURES
ON THE COMPUTATION OF VOWEL HARMONY

By

EILEEN BLUM

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfilment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Linguistics

Written under the direction of

Adam Jardine

And approved by

New Brunswick, New Jersey

January 2023

ABSTRACT OF THE DISSERTATION

The effects of non-linear data structures on the computation of vowel harmony

by EILEEN BLUM

Dissertation Director:
Adam Jardine

This dissertation develops a new theory of autosegmental locality for vowel harmony patterns. Vowel harmony (vh) is a type of pattern in which vowels within a word assimilate to a particular subsegmental feature. Phonological theory has proposed a variety of representational structures to describe the relationships between subsegmental features but little is known about the computational effects of these structures. In this dissertation I use Formal Language Theory, a subfield of computer science, to compare the computational complexity of vh patterns represented over strings and multi-tiered autosegmental representations (ARs). This comparison determines that multi-tiered ARs with “bottle brush” structures (Clements 1976; Hayes 1990; McCarthy 1988; Padgett 1995) are preferable to strings because they reduce the complexity of vh and create more concise descriptions of vh patterns. I extend Jardine (2016b)’s and Jardine (2017a)’s Autosegmental Strictly Local (ASL) to a new complexity class called ASL^{VH} which encompasses vh patterns that are local over multi-tiered ARs. This new class crosscuts the established subregular stringset hierarchy (Heinz 2018; Heinz, Rawal, and Tanner 2011; Rogers et al. 2013; Rogers and Pullum 2011) because it includes patterns which when represented over strings are strictly local like in Akan, Bayinna Orochen, and Kinande; strictly piecewise like in Finnish; locally testable like in Tutrugbu; and it excludes the unattested first-last-harmony pattern which is star-free (Lai 2015; Jardine 2019). The ASL^{VH} class encompasses vh patterns with both opaque and transparent vowels and predicts a new restriction on the locality of transparency. A contrast in the harmonic feature is shown to have no effect on the complexity of opaque vowels but it vastly increases the complexity of patterns with transparent vowels like the one in Eastern Meadow Mari.

Acknowledgements

First, I must express my sincere gratitude to my advisor Adam Jardine. The first graduate class he taught as a Rutgers professor inspired my dissertation topic. Adam has been patient with me through all of my life struggles but always believed in me and this work even when I wasn't so sure.

I also want to thank the other members of my committee for their contributions to my graduate career. Akinbiyi Akinlabi supported my initial experimental project and ensured that I maintained a careful approach to analyzing African languages on this project. Unfortunately, I ran out of time to include all the languages Akin wanted to see but I appreciate that he always asked the difficult big picture questions. Bruce Tesar has also seen me through my graduate research career. Bruce's meticulous insights have ensured I understand the details of computational work and how it fits within the field of phonology. Thank you also to my external committee member Eric Baković for always encouraging me and ensuring that I accurately represent the patterns which you analyzed 23 years ago. I followed in Eric's academic footsteps getting my BA at UC Santa Cruz then my PhD at Rutgers and I truly appreciate his support which even went as far as flying me out to present my work in southern California.

Next, I must express my thanks to other members of the Rutgers Linguistics department. Huge thanks to Paul de Lacy who advised and encouraged my first research project in experimental phonology. Paul always stood up for his students and taught me valuable skills to cope with the struggles of being an academic. Thank you to Veneeta Dayal who forgave my transgressions in semantics and gave me the credits I needed anyway. Thank you also to my fellow graduate students Shu-hao Shih and Christopher Oakden who always had a reassuring response for my anxious questions and provided peer support and friendship at conferences. I must express both gratitude and love for Hazel Mitchley and Livia Camargo-Souza who were the best roommates for six and five years, respectively; they were always there to discuss linguistics and/or emotions at home and encouraged me to remain social as much as possible. Thanks also to the other members of my cohort Lydia Newkirk and Ang Li who tried to help me understand semantics and shared their stories. And I want to thank Natasha Chemey for being an amazing friend to me and always there when I needed to talk about anything.

Thank you especially to Rutgers University for providing excellent health insurance which allowed me to get the help I needed on my mental health journey.

I must also thank the professors and friends I made at UC Santa Cruz who inspired and encouraged me to pursue my PhD in the first place. Junko Ito and Jaye Padgett who taught my undergraduate phonology courses and helped me compile a writing sample which turned into my first research project. Jim McCloskey who inspired my love of the Irish language which eventually became the focus of my experimental research and fieldwork. Donka Farkas who recognized my enthusiasm for linguistics and provided a list of schools I should apply to, including Rutgers. And my closest friends Maxie Corbin and Will Buchannan without whom I may

not have survived Jorge Hankamer and Amy Rose-Deal's Syntax 1 class.

Lastly, I must thank my family for all of their support. My parents have supported me in multiple ways as I navigate life and I doubt I would be as healthy and happy as I am today without their help. My mom Golda Blum always has the best financial advice and my dad Eric Blum even went as far as to proofread this entire dissertation. My cousins Matthew and Amy Marks who included me in their family Thanksgiving when my own family was 3000 miles away. Sue Jagoda and my bff Ally Jagoda who have always welcomed me with open arms and been my second family during many holidays. Last but not in any way the least, this dissertation is dedicated to my feline co-author Gina. She patiently taught me how to speak her dialect of cat and she keeps me on a strict schedule which is extremely valuable for my productivity. Gina found me as I was starting this dissertation and has provided much needed emotional support throughout the years it has taken me to write it.

Contents

Abstract	ii
Acknowledgements	iii
1 Introduction	1
1.1 What is vowel harmony?	2
1.2 Representational Structures	5
1.3 Overview of arguments	6
1.3.1 Phonotactics	6
1.3.2 ARs	7
1.4 Predictions	10
1.5 Dissertation Outline	11
2 The status of autosegmental representations in analyses of vowel harmony	13
2.1 Previous Generative Theories	13
2.2 Formal Language Theory	19
2.3 Conclusion	23
3 Formal introduction	24
3.1 The Subregular Hierarchy	24
3.1.1 Strictly Local	26
3.1.2 Strictly Piecewise	28
3.1.3 Locally Testable	30
3.1.4 Tier-based Strictly Local	31
3.1.5 Autosegmental Strictly Local	33
3.1.6 Types of tiers	35
3.2 Autosegmental Representations (ARs)	35
3.2.1 Multi-tiered ARs	36
3.2.2 Representational assumptions	37
3.3 Definition of Constraints	39
3.4 Summary	44

4	Neutral vowels are simple	46
4.1	Blocking vowels: Akan	48
4.1.1	Blocking over Multi-tiered ARs	50
4.1.2	Surface diffusion is local	53
4.1.3	Blocking over strings	53
4.2	Transparent vowels: Finnish	55
4.2.1	Transparency over multi-tiered ARs	57
4.2.2	Surface iteration can be local	60
4.2.3	Transparency over strings	61
4.3	Conclusion	64
5	ARs reduce the complexity of neutral vowels	66
5.1	Tutrugbu	66
5.1.1	ATR Harmony	68
5.1.2	Unbounded circumambience (UCA) over ARs	71
5.1.3	UCA over strings	76
5.2	Eastern Meadow Mari	79
5.2.1	EMM transparency over multi-tiered ARs	82
5.2.2	Implications of complexity	84
5.2.3	EMM over strings of vowels	86
5.3	Conclusion	88
6	ARs provide concise descriptions of patterns	90
6.1	Baiyinna Orochen	90
6.1.1	[+long] vowels	92
6.1.2	[+high] vowels	94
6.1.3	B. Orochen over multi-tiered ARs	97
6.1.4	B. Orochen over strings	103
6.2	Kinande	108
6.2.1	Kinande over multi-tiered ARs	110
6.2.2	Kinande over strings	116
6.3	Conclusion	119
7	Conclusion	120

List of Acronyms

AR	Autosegmental Representation
ASL	Autosegmental Strictly Local
ATR	Advanced Tongue Root
CNL	Conjunction of Negative Literals
EMM	Eastern Meadow Mari
FLH	First-Last Harmony
FO	First Order logic
FS	Full Specification
FSC	Forbidden Substructure Constraint
FSG	Forbidden Substructure Grammar
IPA	International Phonetic Alphabet
LT	Locally Testable
LTI	Local Test Invariance
LTT	Locally Threshold Testable
MSO	Monadic Second Order logic
NC	Non-Counting
NCC	No Crossing Constraint
NL	Negative Literal
OCP	Obligatory Contour Principle
ODT	Optimal Domains Theory
PT	Piecewise Testable
SC	Subsequence Closure
SF	Star Free
SL	Strictly Local
SP	Strictly Piecewise
SSC	Suffix Substitution Closure
TBU	Tone-Bearing Unit
TSL	Tier-based Strictly Local
UCA	Unbounded Circumambient
WSA	Wide Scope Alignment

1 Introduction

The main hypothesis of this dissertation is that vowel harmony patterns in natural language can be described using a set of local constraints over multi-tiered autosegmental representations (ARs). The overall framework adopted here uses a computational perspective to evaluate the complexity of patterns. Previous work within this framework has studied the expressive power of two-tiered ARs to represent tone patterns (Jardine 2016b, 2017a, 2017b, 2019). I follow this work by investigating the expressive power of ARs with more than two tiers which represent the variety of subsegmental features that can affect vowel harmony patterns. I utilize methods from Formal Language Theory (FLT), which is a field of computer science that studies the complexity of patterns. Using these methods, I compare the computational complexity of a variety of vowel harmony patterns when they are represented over two different abstract data structures: strings and ARs.

My comparison reveals that ARs provide a theory of vowel harmony surface well-formedness that is preferable to the existing theories using strings because:

- (a) ARs reduce the computational complexity of non-local vowel harmony patterns as compared with strings; and
- (b) ARs provide more concise descriptions of local vowel harmony patterns than strings do.

Locality here refers to the possible structure of the constraints: local constraints can only forbid pieces of larger structures; and those pieces are called substructures. Both of these arguments in favor of representing vowel harmony with ARs lead me to propose a new computational complexity class called Autosegmental Strictly Local (ASL^{VH}). This new class is defined as including only vowel harmony patterns which can be described locally using a restricted logic called conjunctions of negative literals (CNLs) that only can conjoin a series of negated substructures. A complexity class is a concept from FLT which allows us to categorize patterns according to their computational properties. These complexity classes are mathematically defined and helpful for determining the expressivity of data structures based on the expressivity of the logic needed to describe patterns given different kinds of structures. Phonologists have applied these methods to the study of phonological patterns and a significant body of work in computational phonology shows that phonological generalizations are properly contained within the regular class of patterns represented over strings (Heinz and Idsardi 2013). Additional work has further established a subregular hierarchy of classes for patterns represented over strings (Heinz, Rawal, and Tanner 2011; Rogers and Pullum 2011; Rogers et al. 2013). I show that the new ASL^{VH} class cross-cuts this subregular hierarchy because attested vowel harmony patterns which fit into different classes over strings can be described locally using a CNL over multi-tiered ARs.

1.1 What is vowel harmony?

Vowel harmony is traditionally considered an assimilatory process that changes the underlying form of a word into a surface form with a restricted set of vowel feature combinations. Assimilation describes a type of input-output mapping in which an input form contains one segment that is associated to a feature and that feature is then shared with, copied, or spread onto other segments. Vowel harmony includes a variety of patterns in which vowels assimilate in a specific feature. For example, a language might have patterns in which vowels assimilate to a round, advanced tongue root (ATR), or height feature. Natural languages can also combine more than one vowel harmony pattern. In transformational analyses the vowels that undergo assimilation are called targets and the one that they assimilate to is called the trigger. In the case of vowel harmony the trigger and target segments are only vowels and assimilation ignores all consonants. So the output surface AR of a word that has undergone vowel harmony assimilation contains more than one vowel with the same feature as the trigger. I use “word” here to mean the phonological word and throughout the dissertation a word consists of the phonological domain of vowel harmony.

In phonology, the set of restrictions on the possible sound combinations that can be found in surface forms is called phonotactics (Chomsky and Halle 1965). In phonotactic terms vowel harmony is a pattern in which words can only contain certain combinations of vowels. Vowel combinations are restricted by the presence of specific vowel features. As Clements (1976)’s states, “All vowels in a word must be drawn from one or the other of two mutually exclusive sets.” For example, Akan utilizes ATR harmony and so an Akan word with only [-low] vowels can contain [+ATR] vowels or [-ATR] vowels **but not both**. A single vowel cannot be both [+ATR] and [-ATR] at the same time because of the physical mechanism that the ATR feature represents: a tongue root cannot be both advanced and retracted at the same time. Except where otherwise noted, the Akan examples presented in this dissertation are adapted from Clements (1976) and the harmony pattern is supported by additional data from Dolphyne (1988).

- (1) Akan words with only [-low] vowels

<u>[+ATR]</u>	<u>[-ATR]</u>
a. tie ‘listen’	d. ɔbɛjɛr ‘he came and did it’
b. obejii ‘he came and removed it’	e. wɔbɛnɔm? ‘you will drink it’
c. wubenum? ‘you will suck it’	

The Akan words with only [-low] vowels in (1) clearly exemplify exactly what Clements (1976) described. The words in (1a-c) contain only [+ATR] vowels and the words in (1d-e) contain only [-ATR] vowels. In short Akan presents a classic example of ATR harmony which can be seen by looking only at attested surface forms.

We don't have to know anything about how these words acquired their final surface forms to know that Akan uses ATR harmony.

In addition, there are some vowels which do not obey the basic vowel harmony generalization described above but are found in languages which do otherwise obey that generalization. Such vowels can either be blocking or transparent to the assimilation of a vowel feature. On the surface blocking means that the vowels on either side of a blocking vowel have different harmonic feature values. Surface transparency means that vowels on either side have the same harmonic feature value which differs from the harmonic feature value of the transparent vowel(s). When grouped together, transparent and blocking vowels are referred to as neutral vowels (Akinlabi 2009; Krämer 2001; Pulleyblank 1996; van der Hulst and Smith 1986; van der Hulst 2016). As you may recall, Akan exemplifies a basic ATR harmony pattern in words with only [-low] vowels. This is because [+low] vowels in Akan are blockers. Following O'Keefe (2004) and Casali (2012) the [+ATR, +low] vowel in the right column of (2) is [æ] rather than the [ɜ] used in Clements (1976).

(2) Akan words with [+low] blockers

<u>[+ATR, +low] [æ]</u>	<u>[-ATR, +low] [a]</u>
a. pɪræko 'pig'	c. obisari 'he asked'
b. mɪkəkæri 'I go and weight it'	d. okog ^w ari [?] 'he goes and washes'

The blocking effect of the [+low] vowels in (2) can be described as follows: the vowels on either side of a [+low] vowel or sequence of [+low] vowels have different ATR feature values within the same word. So while the words in (1) could only contain vowels that are either [+ATR] or [-ATR], including a [+low] vowel allows the words in (2) to contain both [+ATR] and [-ATR] vowels.

Transparent vowels have the opposite effect in that they require vowels on either side of a span of transparent vowels to have the same value for the harmonic feature. For example in Finnish, words without [-back,-round,-low] vowels can have either [+back] vowels or [-back] vowels, **but not both**. A single vowel cannot be both [+back] and [-back] at the same time because it is not possible to raise both the back and the front or center of the tongue at the same time. The Finnish examples presented in this dissertation are copied from Ringen and Heinamaki (1999), notated as [RH] and Välimaa-Blum (1986), notated as [VB].

(3) Finnish words with only harmonizing vowels

<u>[-back]</u>	<u>[+back]</u>
a. pøytæ ‘table’ [RH]	d. pouta ‘fine weather’ [RH]
b. kæntæ: ‘turn’ [VB]	e. murta: ‘break’ [VB]
c. tykætæ ‘like’ [VB]	f. kokata ‘cook’ [VB]

The Finnish words in (3) exemplify a basic back harmony generalization. The words in (3a-c) contain only [-back] vowels and the words in (3d-f) contain only [+back] vowels. Again it is clear from looking only at the attested surface forms that Finnish uses back harmony. Similar to Akan, Finnish also includes words that contain both [+back] and [-back] vowels, but only if the [-back] vowels are also [-round] and [-low]. Unlike Akan, the vowels on one side of a span of these [-back,-round,-low] vowels must be associated to a feature with the same value as the vowels on the other side of that span of [-back,-round,-low] vowels. These Finnish [-back,-round,-low] vowels [i] and [e] are thus called transparent because back harmony appears to have skipped over them. All of the examples below are taken from Välimaa-Blum (1986).

(4) Finnish back harmony skips over transparent vowels [-back,-round,-low] [i, e]

<u>[-back]</u>	<u>[+back]</u>
a. æitiæ ‘mother’	d. ruveta ‘start’
b. ky:neltæ ‘from the nail’	e. tuolia ‘chair’
c. værikæs ‘colorful’	f. lukea ‘read (inf.)’

In (4a-c) the vowels on either side of [i] and [e] are only [-back] and in (4d-f) the vowels on either side of [i] and [e] are only [+back]. So (4d-f) show how transparent vowels can allow Finnish words to contain both [+back] and [-back] vowels, but only if the [-back] vowels are also [-round,-low].

In this dissertation I analyze a variety of vowel harmony patterns including those with blocking and those with transparent vowels. I compare their abstract representations over different data structures in order to use the vowel harmony patterns as evidence for the hypothesis that vowel harmony can be represented locally using ARs and is thus ASL^{VH} . In the following subsection I introduce the data structures used throughout this dissertation.

1.2 Representational Structures

The advantages of ARs mentioned above result from their expanded string-like structure. Strings consist of a series of elements that are connected by a unary ordering relation called successor which indicates direct adjacency from left to right and is represented with a rightward arrow, as shown on the left side of (5). ARs consist of multiple strings called tiers and elements on different tiers are connected by a binary association relation. Association is represented with a straight line, as shown on the right side of (5).

(5) Representational data structures

Strings

a → b → c
or 1 → 2 → 3

Autosegmental representations (ARs)

a → b → c or a → b
| | | | |
1 → 2 → 3 1 → 2 → 3

Multi-tiered ARs are an expansion of the two-tiered ARs originally proposed by Goldsmith (1976) for representing phonological tone and adapted by Clements (1976) for representing vowel harmony. Multi-tiered ARs contain more than two tier of strings with elements on each tier connected by a unary successor relation and elements on different tiers connected by a binary association relation, as shown in (6) below.

(6) Multi-tiered AR with vowels and features

+F → -F
| | |
V → V → V
| | |
+G → -G → +G

I adopt a “bottlebrush” feature representation in which each subsegmental feature is represented on a separate tier and is associated directly to a vowel (Clements 1976; Hayes 1990; McCarthy 1988; Padgett 1995); so each vowel is associated to multiple featural tiers. In addition, I use binary features so vowel features like [± back], [± high], [± round], etc. are represented on separate tiers and each is associated to at least one vowel.

These ARs are derived by concatenation and obey restrictions on the structures they can possibly represent. Concatenation is a mathematical operation on strings which Jardine and Heinz (2015) adapted to derive ARs in the same way that strings are derived. Essentially each vowel and its singular feature associations make up a single “primitive” which is concatenated to adjacent primitives. The concatenation operation thus creates a full AR by joining individual vowel/feature primitives to create an ordered string of vowels which are associated to their respective subsegmental features. The Obligatory Contour Principle (OCP) prevents features with identical values from being adjacent on a tier so they are merged into one multiply associated feature. The

concatenation of adjacent primitives also prevents violations of the No Crossing Constraint (NCC) and can ensure Full Specification (FS) of an AR. In this way the ARs I use obey the standard assumptions in the field of autosegmental phonology: the NCC, and the OCP. In addition, these ARs are able to represent vowel harmony without relying on underspecification. Further details of these restrictions and the derivation of ARs via concatenation are provided in chapter 3. In short, while multi-tiered ARs expand on the structures provided by strings and two-tiered ARs they are still sufficiently restrictive to be useful for defining a theory of locality in natural language phonology.


1.3 Overview of arguments

1.3.1 Phonotactics

In this dissertation I investigate the possible restricted sets of vowel feature combinations that can be found in the surface forms of words. By focusing only on the surface forms of words I am able to determine the computational complexity of the forms which result from a vowel harmony assimilation process. It is useful to fully understand the computational properties of possible output surface forms before one tries to determine the computational properties of the input-output map. Focusing only on the set of possible output ARs allows for a deeper investigation into the properties of ARs and the effects they can have on the computation of vowel harmony patterns. For example my formal investigation into ARs helps to clarify the details of a variety of vowel harmony patterns and reveals the simplicity of patterns that were previously considered complex.

The phonotactic constraints used to describe vowel harmony patterns throughout this dissertation have a few different names due to their their logical function and structure. As with most phonotactic markedness constraints they are negative because they can only forbid something; they cannot require something. In mathematical logic a statement which consists solely of negation and some piece of structure (a substructure) is called a negative literal (NL). A negative literal has the following structure: $\neg w_0$ where w_0 represents some substructure. The particular type of negative literal used throughout this dissertation is called a Forbidden Substructure Constraint (FSC) (Heinz, Rawal, and Tanner 2011; Jardine 2017b; Rogers et al. 2013). FSCs forbid a piece of an abstract phonological representation from occurring within a full grammatical representation—much like phonotactic markedness constraints.

(7) Forbidden Substructure Constraint (FSCs)

over strings $\neg \text{oo}$ over ARs

FSCs can be written using logical negation (\neg) or the linguistic symbol for ungrammatical, an asterisk (*).

1.3.2 ARs

Vowel harmony patterns are describable over both strings and ARs but the relative independence of subsegmental features within vowel harmony patterns motivates their representation with ARs. Harkening back to Clements (1976) who first adapted Goldsmith (1976)'s autosegmental representations for vowel harmony, the pattern of subsegmental feature assimilation across all vowels in a word is clearly and efficiently represented by separating segments from their subsegmental features. When represented on separate tiers we can clearly see how these features “behave with relative autonomy” (Clements 1976) such as by associating with multiple different vowels, by skipping vowels, and by restricting or interrupting the association of vowels to other features or feature values. In addition, separating features from vowels in ARs clearly shows the differences between the computations of vowel harmony assimilation with different types of neutral vowels. In other words, when we represent them with ARs the behavior of vowel features with respect to vowel harmony and neutral vowels is clarified.

The advantages of ARs outlined in (a) and (b) above make them preferable to strings for a restrictive theory of vowel harmony. Jardine (2016b) also argued that a restrictive theory of phonology prefers to enrich the representation rather than increase the expressivity of a grammar which can cause overgeneration. For example, over strings a long distance vowel harmony pattern like ATR harmony in Tutrugbu (Essegbey 2019; McCollum and Essegbey 2020; McCollum et al. 2020) is more complex than many other patterns and its description requires a logic which can make conditional statements. The Tutrugbu nine-vowel inventory is shown in Table 1 below.

Table 1: Tutrugbu Vowels

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

In Tutrugbu prefixes have the same value for the feature Advanced Tongue Root (ATR) as the root-initial vowel. Some of the examples in (8) below include polysyllabic roots or multiple prefixes so morphemes are separated by a plus symbol (+) and the rightmost morpheme is the root. The nouns below are taken from McCollum et al. (2020) and the verbs are from Essegbey (2019).

(8) Tutrugbu full ATR harmony

Nouns

- a. e+bu ‘CL1+dog’
- b. o+pete ‘CL3+vulture’
- c. ɪ+da ‘CL4+copper’
- d. bu+wi ‘CL8+axe’

Verbs

- e. i+gi+tseɔe ‘I did not tell’
- f. o+bo+nyi ‘You will know’
- g. ɛ+ge+za ‘I did not stay’
- h. ɔ+bo+ba ‘You will come’

Tutrugbu also has a [+low] vowel [a] which can block ATR harmony but only if the word-initial vowel is [+high] and there is a [+ATR] vowel to its right. This pattern is called unbounded circumambient (UCA) because the blocking vowel must be surrounded on both sides by specific other vowels—an initial [+high] vowel on the left and a [+ATR] vowel on the right—but those can be any distance away from the blocker as shown in (9) below.

(9) UCA blocking of [+ATR] harmony (McCollum et al. 2020)

- a. ɪ+ba+wu ‘1sg+Fut+climb’
- b. ɪ+tu+ka+a+ba+ba+wu ‘1sg+Neg+Pfv+Prog+Vent+Vent+climb’

The logical description of Tutrugbu ATR harmony with UCA blocking over strings of vowels requires a conditional statement. The use of a conditional in a logical statement means that the level of logic needed to describe Tutrugbu over strings of vowels is relatively high so the UCA blocking pattern is not local on the surface. However, expanding the representation to include multiple tiers of strings connected by association makes the UCA blocking pattern describable in a local way using the lowest level of logic: a conjunction of negative literals (CNL).

Even with a CNL over strings the descriptions of local vowel harmony patterns like round harmony in Baiyinna Orochen (B. Orochen) (Li 1996; Walker 2014a) can be very long and difficult to read. The vowel inventory of B. Orochen is shown in Table 2 below.

Table 2: Baiyinna Orochen Vowels

	-back	+back	
+high	i, i:	u, u:	+ATR
	ɪ, ɪ:	ʊ, ʊ:	-ATR
-high	iẽ	ə, ə:	+ATR
	iẽ	a, a:	-ATR
	-round	+round	

In B. Orochen all vowels in a word must have the [+round] feature if the leftmost vowel is [+round, -high, -long] [o] or [ɔ]. Without an initial [o] or [ɔ] and if the initial [+round] vowel is also [+long] the remaining vowels are [-round]. This pattern holds both within roots and across morpheme boundaries to suffixes as shown in (10). Here the leftmost morpheme is the root and the same plus symbol (⊕) is used to represent morpheme boundaries. All the B. Orochen data in this dissertation comes from Li (1996).

(10) Round harmony

[+round]	[-round]
a. tʃolopon ‘morning star’	e. targan ‘field, garden’
b. ɔɾkto ‘hay’	f. ko:məxə ‘windpipe’
c. somsoʃ⊕jo ‘pasture (indef.acc)’	g. faʃə⊕lɛ:⊕tʃə ‘to launch’
d. ɔlɔ⊕jo ‘fish (indef.acc)’	h. ɡɔ:l⊕ja ‘policy (indef.acc)’

The B. Orochen [+round] harmony pattern is also blocked by [+high] vowels. Unlike Tutrugbu, B. Orochen exemplifies a standard case of blocking with no special restrictions. Some examples are shown in (11) below.

(11) [+high] vowels block [+round] harmony

a. owon⊕dulə ‘pancake (destin)’	d. ɔɾn⊕dula: ‘reindeer (destin)’
b. bolboxi⊕wə ‘wild duck (def.acc)’	e. tʃəlɪk⊕pa ‘cloud-shaped design(def.acc)’
c. moliktə ‘a kind of wild fruit’	

The examples in (11) clearly show that a B. Orochen word can contain both [+round] and [-round] vowels when they are separated by a [+high] vowel. Over strings of vowels the B. Orochen round harmony pattern can be described with a CNL but it requires 56 constraints. However, those descriptions are condensed and the

number of constraints significantly reduced when the representation is expanded; B. Orochen is describable with only three constraints over ARs. Thus by using ARs I can describe a wide variety of vowel harmony patterns from languages spoken around the world in a local way that is relatively short and easy to parse.

1.4 Predictions


Together the ASL^{VH} class I propose with the ARs I use make specific predictions about the expressivity of vowel harmony. First, the ASL^{VH} class predicts that vowel harmony patterns with blocking are local over ARs. This prediction is supported by evidence from two languages with blocking patterns that fit into different complexity classes over strings. The complexity classes on the established subregular hierarchy for strings are mathematically defined by the expressivity of the logic needed to describe the patterns within them. One of the simplest classes is the Strictly Local (SL) class (Heinz 2018; McNaughton and Papert 1971; Rogers et al. 2013). SL contains patterns that are describable with a CNL over strings that are ordered with strict adjacency—also called successor—like ATR harmony with blocking in Akan. A class which is strictly more complex than SL is the Locally Testable (LT) class (Heinz 2018; Rogers et al. 2013) which contains patterns that require a more expressive logic called propositional logic to describe them such as the ATR harmony with UCA blocking in Tutrugbu. However, both the Akan and Tutrugbu patterns are local over ARs. In addition, over ARs the Akan pattern is described with a CNL of only three FSCs and the Tutrugbu pattern is described with a CNL of only four FSCs.

Second, the ASL^{VH} class predicts that transparent vowels do not contrast in the harmonic feature and as such transparency can be represented locally on the surface without relying on underspecification. This prediction is supported by evidence from two languages with transparency patterns that fit into different complexity classes over strings. Another one of the simplest classes of string patterns is called Strictly Piecewise (SP) (Heinz 2018; Rogers et al. 2010, 2013). SP contains patterns that are describable with a CNL over strings that are ordered with the less strict precedence like backness harmony with transparency in Finnish. While backness harmony with positional transparency in Eastern Meadow Mari (EMM) is LT only the Finnish pattern is local over ARs. The Finnish pattern is described with a CNL of only four FSCs over ARs and illustrates that transparency with no harmonic contrast can be represented locally without relying on underspecification. The EMM pattern, however, is excluded from the ASL^{VH} class because it utilizes two transparent vowels which contrast in the harmonic feature, $[\pm\text{back}]$.

Lastly, the ASL^{VH} class excludes the unattested pattern called First-Last Harmony (FLH) (Heinz 2018; Lai 2015). As a vowel harmony pattern FLH is one in which only the first and last vowel of any word must have the same feature value; any vowels can intervene between the first and last one with no effect and are

essentially treated as transparent.

(12) First-Last Harmony is non-local



Lai (2015) claimed that first-last vowel harmony patterns are unattested in natural language and Walker (2011) first analyzed EMM as FLH which chapter 5 shows is not true.¹ FLH provides a well-documented example of why harmonic contrast among transparent vowels cannot be represented locally over ARs. Using binary features predicts that a grammatical word could exist in which the first and last vowel are associated to separate iterations of a feature with the same value but the intervening vowels are associated to alternating feature values, as shown in (12) below. As the number of intervening vowels grows so too does the length of the string on the harmonic F feature tier. In this case the distance between the first and last F feature depends upon the length of the vowel string. This dependency makes it impossible to describe FLH using only a CNL over ARs; so FLH is a long-distance pattern because it cannot be local over ARs and thus is not ASL^{VH} .

1.5 Dissertation Outline

The remainder of this dissertation has the following structure. Chapter 2 outlines the theoretical history of vowel harmony analyses as motivation for a study of ARs. The second chapter explains that ARs are useful for representing vowel harmony patterns because they are sufficiently expressive to describe such patterns and their ability to reduce the complexity of patterns in general helps a theory of vowel harmony to remain maximally restrictive.

Chapter 3 provides the formal background to define the ARs I use and the new ASL^{VH} class I propose. I further explain the hierarchy of complexity classes used to categorize natural language phonology and propose the ASL^{VH} hypothesis: vowel harmony patterns are describable using a CNL of phonotactic restrictions over ARs and are thus ASL^{VH} .

Chapters 4, 5, and 6 provide empirical evidence for my claims about the computation of vowel harmony. Chapter 4 defines neutral vowels and investigates the effects they can have on the computation of vowel harmony with analyses of blocking vowels in Akan and transparent vowels in Finnish. The fourth chapter

¹C'Lela includes a height harmony pattern which has also been considered as a potential attestation of FLH but Lai (2015) points out that the data does not fully support that analysis.

illustrates the structural differences between blocking and transparency and argues that these differences cause vowel harmony patterns with these types of neutral vowels to fall into different complexity classes over strings but not over ARs. In addition, the fourth chapter argues that Finnish transparency is local even without the use of underspecification on the surface.

Chapter 5 includes analyses of two complex patterns with neutral vowels and discussion of how ARs can reduce the complexity of Unbounded Circumambient (UCA) blocking in Tutrugbu but not positional transparency in Eastern Meadow Mari (EMM). Building on the previous chapter, the fifth chapter argues that blocking is inherently local regardless of whether the blocking vowels contrast in the harmonic feature or not. On the other hand, transparency is local with no harmonic contrast but the complexity of patterns with transparency is increased when transparent vowels contrast in the harmonic feature.

Chapter 6 argues that ARs vastly reduce the number of constraints needed to describe vowel harmony patterns that are already local over strings. Evidence is shown with analyses of round harmony in B. Orochen and ATR harmony in Kinande.

Chapter 7 concludes.

2 The status of autosegmental representations in analyses of vowel harmony

In this chapter I outline the theoretical motivations for studying vowel harmony over autosegmental representations (ARs). At a basic level ARs are a data structure consisting of more than one string of elements with connections between the elements on each string and between elements on different strings. Each string in an AR is called a tier; the unary ordering relation which connects elements on a tier is called successor; and the binary relation which connects elements on the vowel tier and on each feature tier to one another is called association.

I investigate the computational complexity of vowel harmony as a surface phonotactic restriction rather than a process that changes an input into a particular output form. Phonotactic restrictions describe the possible forms that can be found on the surface regardless of the input structure. A phonotactic account of vowel harmony thus describes only the vowel and feature combinations that are possible on the surface. Any mention of triggers or targets is used only in reference to external input-output mapping accounts of the patterns I analyze.

This dissertation constitutes the first in-depth formal language theoretic study of vowel harmony over multi-tiered ARs. Every vowel harmony pattern I analyze has previously been described as a transformation, but phonotactics allow for a strong focus on the computational properties of representations. My formal investigation into multi-tiered ARs helps to clarify the details of vowel harmony and reveals the simplicity of patterns that were previously considered complex. It is also useful to fully understand the computational properties of possible output surface forms before one tries to determine the computational properties of the input-output map. Focusing only on the set of possible output ARs allows for a deeper investigation into the properties of ARs and the effects they have on the computation of vowel harmony patterns. In this chapter I situate ARs within the context of the input-output mapping theories of vowel harmony that have often been cited in the literature and present theoretical motivations for using them.

2.1 Previous Generative Theories

Early work identified two major types of vowel harmony transformations: Dominant-Recessive and Root/Stem-controlled (Levergood 1984; Ringen 1975). Dominant-recessive vowel harmony is a pattern in which triggering vowels are dominant in some way—often due to external properties, such as prosodic prominence, strong position within the word, etc.²—and targets are recessive. Root/Stem-controlled vowel

²See Walker (2011) for discussion of positions of phonological prominence relevant for vowel harmony.

harmony is a pattern in which a vowel within a root or stem triggers harmony and vowels within affixes are targets. Derivational rule-based analyses such as Ringen (1975) and Levergood (1984) assume that a vowel harmony rule transforms some underspecified underlying form into a fully specified surface form. For example the derivations of two Finnish words in (13) contain the underlying forms of the word shown at the top surrounded by slashes (/ /). The underlying forms contain capital letters to represent vowels which have not yet been specified for backness. The following vowel harmony rule is then applied to these underlying forms in order to spread the back feature from the initial underlyingly specified vowel onto the remaining unspecified vowels: [+syllabic] → [αback] / [+syllabic, αback] (C₀V₀)*__ (Ringen 1975). The results of applying this vowel harmony rule are shown in the middle row of (13).

(13) Derivation using vowel harmony rule (adapted from Ringen (1975))

underlying form:	/pøUtA/	/po UtA/
	↓	↓
vowel harmony:	pøytæ	pouta
	↓	↓
surface form:	[pøytæ]	[pouta]

The surface forms are shown at the bottom surrounded by square brackets ([]). In other words, the trigger is specified for the harmonic feature in the underlying form, but targets are not. A vowel harmony rule then applies to assimilate the unspecified targets to the harmonic feature value of the trigger.


The first autosegmental theory of vowel harmony was published around the same time as the introduction of autosegmental theory into phonology (Clements 1976; Goldsmith 1976). Clements (1976) posited a Well-Formedness Condition which motivates feature spreading in order to ensure that all elements on one tier of an AR are connected via an association relation to some element on another tier of the same AR. The result is Full Specification which means that within an AR all elements on one tier are associated to some element on another tier.³

(14) Well-Formedness Condition (adapted from Clements 1976)

- a. all vowels are associated to at least one feature on each tier
- b. all features on each tier are associated to at least one vowel
- c. for each feature tier, association lines never cross


³An in-depth discussion of the autosegmental structures and restrictions I utilize throughout this dissertation is found in the next chapter.

d. Example: [obejii] ‘he came and removed it’ (Akan)




The Akan word in (14d) illustrates that obeying Clements (1976)’s Well-Formedness Condition results in a fully specified AR. Each vowel is associated to a feature on each of the ATR and low tiers. This is only possible because the single [+ATR] feature is multiply associated to all of the vowels and the single [-low] feature is multiply associated to all the vowels. Neither ATR nor low association lines are crossed. Multiple association thus allows the AR in (14d) to be fully specified and obey the Well-Formedness Condition. In an input-output mapping theory multiple association is said to result from the spreading of a feature’s association from the underlyingly specified vowel onto all the unspecified vowels, as in (15).

(15) Spreading satisfies the Well-Formedness Condition



So the process of feature spreading takes an underspecified AR and creates a surface AR which obeys the Well-Formedness Condition. The Well-Formedness Condition crucially allows blocking to be represented autosegmentally. Any underlyingly unspecified target vowels are associated to the same ATR feature as the trigger but the unspecified vowels preceding the [+low] blocker are not, as shown in (16) below.

(16) [pɪræko] ‘pig’ (adapted from Clements 1976)



The initial vowel in (16) is associated to [-ATR] by default. The [+low] vowel is underlyingly specified as [+ATR] so it is not a target of assimilation and it blocks the spread of [+ATR]. Clements (1976)’s Well-Formedness Condition helped to define ARs and how they could be used to analyze vowel harmony patterns for the next decade and a half and it continues to influence the use of ARs in phonological theory to this day.

For just over a decade after the introduction of autosegmental theory, phonological study included a major emphasis on representations but eventually phonologists' focus changed. Investigations into phonological representations and ARs in particular were deemed valuable under the assumption that "if the representations are right then the rules will follow" (McCarthy 1988). However, the introduction of Optimality Theory (OT) (Prince and Smolensky 1993) popularized a new form of transformational analysis which downplayed the importance of representations and suggested that phonological generalizations fall out from the ranking of universal constraints instead. Baković (2000) showed that vowel harmony can be understood in the OT framework without needing ARs, but he abstracted away from transparency which requires some additional theoretical machinery in OT. Below I discuss two OT analyses of vowel harmony which incorporate subsegmental features into their constraints but shy away from ARs.

Cole and Kisseberth (1994) developed their Optimal Domains Theory (ODT), which suggested that vowels within a word are parsed into phonological domains that contain the trigger and targets of harmony. A word can also contain vowels outside of that domain, which are not targets of harmony. They posited three OT constraints to account for vowel harmony: Wide Scope Alignment (WSA), Expression, and *Insert [F]. WSA constraints state that the left or right edge of a featural domain must be aligned with the left or right edge of either a prosodic or morphological boundary; Expression states that the relevant feature must be affiliated with every anchor (vowel) within the harmony domain; and *Insert[F] prohibits the insertion of a feature. In languages with vowel harmony, *Insert[F] must be ranked below both Expression and the relevant WSA constraint: Expression, WSA \gg *Insert[F]. In Cole and Kisseberth (1994) transparency and blocking result from the crucial ranking of another type of constraint they call "feature occurrence constraints" with respect to Expression and WSA. F-Occurrence constraints restrict the occurrence of a feature for some vowels. The overall rankings for vowel harmony, blocking, and transparency are listed below.

- (17) Rankings from Cole and Kisseberth (1994)
- a. Harmony: Expression, WSA \gg *Insert[F]
 - b. Blocking: F-Occurrence, Expression \gg WSA
 - c. Transparency: F-Occurrence, WSA \gg Expression

These rankings illustrate the three basic generalizations of vowel harmony patterns. The general concept of vowel harmony in ODT is that all vowels within some domain share a feature, as in (17a). In some languages, however, there are vowels which block the spread of a feature so only a subset of successive vowels within the domain are affiliated with that feature, as in (17b). In other languages there are vowels that are not affiliated with the relevant feature but vowels on either side of them are and so left or right edge alignment of a harmony domain can be maintained as in (17c).

Later, McCarthy (2004) posited a similar analysis of nasal harmony in OT. Essentially, segments are parsed into phonological spans (domains) within which the feature value of a trigger spreads onto all other segments. He posited that harmony results from the crucial ranking between a faithfulness constraint and a markedness constraint. The faithfulness constraint was called $FthHdSp(\alpha F)$ and requires an input with a particular feature value to head a span with that same feature value in the output. The markedness constraint was called $*A-Span[F]$ and is violated by adjacent $[F]$ spans. McCarthy (2004) claims that harmony results from the following ranking: $FthHdSp(\alpha F) \gg *A-Span[F] \gg FthHdSp(-\alpha F)$. In addition, directional harmonies result from different rankings of $SpHdL(\pm F)$ and $SpHdR(\pm F)$, which require head segments to be aligned with the left or right edge of a span. He further posits a series of markedness constraints with the form $Head([\beta G, \gamma H, \dots], [\alpha F])$ that require certain types of segments to head spans with a particular feature value, thus acting as feature cooccurrence restrictions. He then claims that “blocking effects result from ranking $*A-Span[F]$ below one of [these] constraints.” The overall rankings from McCarthy (2004) are listed below.

(18) Rankings from McCarthy (2004)

- a. Harmony: $FthHdSp(\alpha F) \gg *A-Span[F] \gg FthHdSp(-\alpha F)$
 - i. left-to-right: $SpHdL(-F) \gg SpHdR(+F)$; $SpHdL(+F) \gg SpHdR(-F)$
 - ii. right-to-left: $SpHdR(-F) \gg SpHdL(+F)$; $SpHdR(+F) \gg SpHdL(-F)$
 - iii. bidirectional: $SpHdL(-F) \gg SpHdR(+F)$; $SpHdR(-F) \gg SpHdL(+F)$
- b. Blocking: $Head([\beta G, \gamma H, \dots], [\alpha F]) \gg *A-Span[F]$

McCarthy (2004)’s theory of headed spans does not account for patterns with transparency like those found in some vowel harmony languages. In order to account for a pattern like Finnish which has two vowels $[i]$ and $[e]$ that are transparent to backness harmony McCarthy (2004)’s theory of headed spans must be supplemented with a constraint like Cole and Kisseberth (1994)’s Expression. Expression would be crucially ranked below a feature cooccurrence constraint, which is non-crucially ranked below $SpHdL[+back]$. This ranking forces the left-to-right spread of a $[+back]$ feature from the trigger onto its targets and prevents the transparent vowel from changing its back feature value.

(19) Example using McCarthy (2004) with transparent [e]

/ruvetæ/	SPHdL[+BK]	SPHdR[-BK]	*[-RND,+BK,-LOW]	EXPRESSION
a. (ru <u>v</u> eta)				*
b. (ru <u>v</u> ata)			*!	
c. (ru <u>v</u> etæ)				**!
d. (ruvet <u>a</u>)	*!			*
e. (ryvet <u>æ</u>)		*!		

The highly ranked SpHdL[+back] constraint in (19) reduces the candidate set to those only with left headed spans of [+back] vowels. As shown in (18ai) SpHdL[+back] is crucially ranked above SpHdR[-back] in order to enforce that harmony applies from left to right. The transparent vowels in Finnish are all [-round, -back, -low] and so the feature cooccurrence constraint, *[-round,+back,-low], prevents the transparent [e] vowel in (19) from changing to its [+back] counterpart via the crucial violation for candidate b. Expression is ranked below these other two constraints because it is violated by both candidates a and c; both candidates violate Expression because the transparent vowel is included within the [+back] span, but is not [+back]. Candidate c is the maximally faithful candidate, which crucially violates Expression a second time because it contains two [-back] vowels. So (19) illustrates that any number of transparent vowels could occur between the two [+back] vowels because a maximally faithful candidate will always accrue one more violation of Expression than the optimal candidate with transparency.

The goal of generative phonology is to develop a theory that is expressive enough to describe how attested natural language sound patterns are computed and restrictive enough not to predict the existence of unattested sound patterns. The OT framework is clearly expressive enough to generate vowel harmony patterns as shown above but the computational power of optimization has also been shown to be too expressive in some cases. For example, the constraints from Cole and Kisseberth (1994) and McCarthy (2004) above can generate harmony, blocking, and transparency patterns; but others have also shown that optimization predicts the existence of patterns we don't see in natural language such as sour grapes and majority rules (Heinz 2018; Heinz and Lai 2013; Lamont 2019; O'Hara and Smith 2019; Smith and O'Hara 2019). One way to understand the restrictions that are needed to prevent a theory from overgenerating phonological patterns is to investigate the computational complexity of those patterns. Computational complexity tells us precisely how expressive a theory must be to describe a pattern and provides categories for grouping patterns according to precise restrictions on their expressivity. In short, understanding the complexity of a pattern tells us exactly how expressive and how restrictive a theory must be to accurately describe that pattern without overgenerating.

2.2 Formal Language Theory

Computer science provides an established method for investigating the computational complexity of patterns in general called Formal Language Theory (FLT). Heinz (2018) claims that the theory of computation on which FLT is based provides a method for measuring the expressivity and restrictiveness of a theory. The theory of computation tells us that patterns can be distinguished based on the expressive power needed to compute them. The Chomsky hierarchy provides a concrete way to quantify these distinctions: patterns are grouped together based on their computational properties and each category of patterns is ranked according to the logical expressivity needed to describe the patterns within it. These categories of patterns are called complexity classes. The complexity classes are mathematically defined which provides a formal metric for both the logical expressivity needed to describe a pattern and the restrictions needed to exclude other patterns from a class. Heinz (2018) argues that OT is expressive enough to compute vowel harmony patterns, but FLT provides a framework for developing a theory of vowel harmony which is also restrictive enough so as not to predict unattested patterns.

In the second decade of the twenty-first century phonologists have used FLT to determine that the computational complexity of phonological patterns including vowel harmony is affected by representational structures (Aksënova and Deshmukh 2018; Chandlee 2014; Chandlee and Heinz 2018; Heinz 2018; Heinz, Rawal, and Tanner 2011). In other words, the complexity of a pattern changes when that pattern is represented over different data structures such as strings or autosegmental representations (ARs). Strings are generated from a finite set of symbols which is called the alphabet and denoted by Σ . A formal language is a set of strings, or stringset, which is defined by certain computational properties. Collectively, formal languages make up a subset of all possible stringsets denoted by Σ^* . While most FLT phonology so far has focused on string patterns, Jardine (2016b) showed that ARs have a major effect on the logical expressivity needed to compute phonological patterns. His work highlighted once again the importance of representations for phonological theory.

Previous string-based analyses of vowel harmony proposed that surface patterns with and without neutral vowels—such as Akan (Casali 2012; Clements 1976; Dolphyne 1988; O’Keefe 2004), Finnish (Nevins 2010; Ringen and Heinamaki 1999; Välimaa-Blum 1986; van der Hulst 2017), and Lokaa (Aksënova 2017; Aksënova and Deshmukh 2018)—generally fit into the Tier-based Strictly Local (TSL) class of patterns represented over strings (Heinz 2018; Heinz, Rawal, and Tanner 2011). TSL patterns are described as local over a string of sounds with a tier projection that contains a subset of the sounds in a word. The specific content of the TSL tier projection is language-dependent. TSL tiers differ from autosegmental tiers in certain crucial ways that will be discussed in detail in the next chapter. For example, Aksënova (2017) claims that ATR harmony in

Lokaa with transparent high vowels is TSL because the pattern can be described as local only on a tier of non-high vowels. The generalization for Lokaa is that non-high vowels have the same ATR feature value and “all high vowels and consonants are transparent for the harmony” (Aksénova 2017). The Lokaa examples below are taken from Akinlabi (2009).

(20) Lokaa (Niger-Congo) (Akinlabi 2009)

- a. èsìsòn ‘smoke’, *è̀sìsòn
- b. è̀sìsòn ‘housefly’, *è̀sìsòn
- c. lè̀jìmà ‘matriclan’, *lè̀jìmà
- d. ékílikà ‘kind of plant’, *é̀kílikà

(21) ATR harmony in Lokaa



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 (21). The Lokaa ATR harmony pattern can be described using a series of negative constraints which forbid adjacent elements with different ATR feature specifications and those constraints are evaluated only over the tier projection of non-high vowels. Each literal within the CNL thus contains two vowels which are adjacent on the tier projection. Because the Lokaa pattern is described by literals with only two vowels it is considered tier-based strictly 2-local (TSL₂). The TSL₂ grammar that generates the Lokaa vowel harmony pattern would look like (22).

(22) TSL grammar for Lokaa harmony

- a. Tier of non-high vowels: T = ε, e, o, ə, ɔ, a
- b. CNL: * $[\alpha\text{ATR}][\beta\text{ATR}]$ = *εε ∧ *εε ∧ *εo ∧ *oε ∧ *εə ∧ *əε ∧ *ɔe ∧ *eɔ ∧ *ɔə ∧ *əɔ ∧ *ɔo ∧ *oo ∧ *aə ∧ *əa ∧ *ao ∧ *oa ∧ *ae ∧ *ea

The TSL grammar in (22) illustrates how 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 which allows the pattern to be described with negative constraints over a subset of the segments in a word.

However, additional vowel harmony patterns have recently been described as even more complex than TSL when represented over strings. Two such patterns are the unbounded circumambient blocking of ATR harmony found in Tutrugbu (McCollum and Essegbe 2020; McCollum et al. 2020) and back harmony with positional transparency found in Eastern Meadow Mari (Vaysman 2009; Walker 2011). The ATR harmony

generalization in Tutrugbu has been described as right-to-left assimilation of prefixes to the ATR feature of the leftmost stem or root vowel. Some examples of full ATR harmony from Essegbey (2019) are shown in (23)-(24) below. I use the large plus symbol (+) as a morpheme boundary; so roots are the rightmost morpheme and each morpheme to the left of a root is a prefix.

(23) Tutrugbu nouns with full ATR harmony

- | <u>[+ATR]</u> | <u>[-ATR]</u> |
|---|--|
| a. e + bu ‘CL1 + dog’ | c. r + da ‘CL4 + copper’ |
| b. i + pete ‘CL4 + vulture’ | d. bu + wi ‘CL8 + axe’ |

(24) Tutrugbu verbs with full ATR harmony

- | <u>[+ATR]</u> | <u>[-ATR]</u> |
|---|---|
| a. i + gi + tse + e ‘I did not tell’ | c. e + ge + za ‘I did not stay’ |
| b. o + bo + nyi ‘You will know’ | d. o + bo + ba ‘You will come’ |

This ATR harmony pattern also includes blocking, but only in a very specific environment. The [+low, -ATR] [a] vowel blocks [+ATR] harmony from the right, but only when the initial vowel is [+high]. This blocking pattern is called unbounded circumambient because the blocker [a] must be surrounded by triggers on each side—[+high] on the left and [+ATR] on the right—and those triggers can be any distance from the blocker. A clear example from McCollum et al. (2020) is shown below in (25) and contrasted with a version of the same verb that does not include blocking in (26).

(25) UCA blocking of [+ATR] harmony

- a. r + tu + ka + wu ‘1sg + Neg + Pfv + climb’

(26) Full [+ATR] harmony, no blocking

- e + ti + ke + e + be + be + wu ‘3sg + NEG + Pfv + Prog + Vent + Vent + climb’

The verb in (25) includes [+high, -ATR] [ɪ] vowels and a [+ATR] [u] vowel on either side of the [+low, -ATR] [a] blocker, but because the initial vowel in (26) is [-high] there is no blocker and all the vowels are [+ATR].

The back harmony generalization in Eastern Meadow Mari (EMM) can be described as applying from left to right triggered by the initial vowel. EMM words that exhibit full backness harmony are shown in (27). All EMM data in this dissertation is adapted from Vaysman (2009).

(27) Eastern Meadow Mari full back harmony

- | | |
|--|--|
| <p>a. Roots</p> <p>i. <i>ɟyɶæ</i> ‘sister’</p> <p>ii. <i>murna</i> ‘tube, pipe’</p> | <p>b. Nom.sg.2p.pl.poss</p> <p>i. <i>tʃødræ</i>+<i>tæ</i> ‘your (pl) forest’</p> <p>ii. <i>kutko</i>+<i>ta</i> ‘your (pl) ant’</p> |
| <p>c. Nom.sg.3p.pl.poss</p> <p>i. <i>ij</i>+<i>næ</i> ‘our year’</p> <p>ii. <i>tam</i>+<i>na</i> ‘our taste’</p> | <p>d. Dative</p> <p>i. <i>ʃør</i>+<i>læn</i> ‘milk (dat.)’</p> <p>ii. <i>kawun</i>+<i>lan</i> ‘pumpkin (dat.)’</p> |

In addition, the two [-high, -low, -round] vowels [-back] [e] and [+back] [ə] behave as transparent only in non-initial positions. This means they can occur interchangeably in the middle and at the end of a word while back harmony skips over them, but a suffix vowel will harmonize with an initial [e] or [ə] as shown in (28).

(28) Eastern Meadow Mari transparency

- | | |
|--|---|
| <p>a. Roots</p> <p>i. <i>teɟgəz</i> ‘sea’</p> <p>ii. <i>jəlme</i> ‘tongue, language’</p> | <p>b. Dative</p> <p>i. <i>imɶə</i>+<i>læn</i> ‘horse (dat.)’</p> <p>ii. <i>ləβe</i>+<i>lan</i> ‘butterfly (dat.)’</p> |
|--|---|

As shown later in this dissertation, both UCA blocking in Tutrugbu and positional transparency in EMM require propositional logic to describe them over string representations. Thus neither are TSL and both are more complex than Akan, Finnish, and Lokaa. Full analyses of both these patterns can be found in chapter 5 along with additional insights.

Increased complexity requires more expressive logics to describe the Tutrugbu and EMM patterns, but relying on more expressive logics also increases the likelihood that a theory of vowel harmony could overgenerate. Jardine (2016b) showed that ARs reduce the complexity of tone patterns when compared with string representations and reduced complexity is appealing for a theory of vowel harmony to be maximally restrictive. This dissertation extends the notion that ARs reduce complexity to investigate the effects of multi-tiered ARs on the computational complexity of a variety of vowel harmony patterns including Tutrugbu and Eastern Meadow Mari. Because these patterns are more complex over strings they provide strong evidence for the usefulness of ARs in a general theory of vowel harmony computation.

Using the FLT framework continues the tradition of exploring the expressivity vs. restrictiveness needed for phonological theories and investigating computational complexity has brought the field of phonology back around to better understand the value of abstract representations.

2.3 Conclusion

In this chapter I have outlined the theoretical history of vowel harmony analyses and explained some of the motivations for focusing on ARs throughout this dissertation. Phonologists have often thought of vowel harmony as an input-output mapping but surface phonotactics provide a strong focus on the computational properties of representations. While the field of phonology has shied away from focusing on representations since the introduction of OT a recent upsurge of work using FLT has reminded us that abstract representations are crucial to understanding the computational properties of phonological patterns. It is also useful to fully understand the computational properties of possible output surface forms before one tries to determine the computational properties of the input-output map because focusing only on the set of possible output ARs allows for a deeper investigation into the properties of ARs and the effects they can have on the computation of vowel harmony patterns. For example my formal investigation into multi-tiered ARs helps to clarify the details of a variety of vowel harmony patterns and reveals the simplicity of patterns that were previously considered complex. In other words ARs are useful for representing vowel harmony patterns because they are sufficiently expressive to describe such patterns and their ability to reduce the complexity of patterns helps a theory of vowel harmony to remain maximally restrictive. Further study of the possible restrictions on ARs could even lead to a theory of vowel harmony which is restrictive enough to avoid predicting unattested patterns.

In this chapter I also explained how vital Clements (1976)'s Well-Formedness Condition has been to the development of ARs over time. The Well-Formedness Condition introduced and motivated principles which are still used to constrain ARs to this day. The multi-tiered ARs I use throughout this dissertation are strongly influenced by the Well-Formedness Condition and in the next chapter I provide all the formal details which show how it is applied.

3 Formal introduction


Formal Language Theory (FLT)—a subfield of theoretical computer science—provides precise mathematical and logical methods for determining the computational expressivity of structures and complexity of patterns. Phonologists have applied these methods to the study of phonological patterns such as consonant and vowel harmony over string representations (Aksënova and Deshmukh 2018; Chandlee 2014; Chandlee and Heinz 2018; Heinz 2018; Heinz, Rawal, and Tanner 2011; McMullin 2016; McMullin and Hansson 2014) as well as tone patterns over two-tiered autosegmental representations (ARs) (Jardine 2016b, 2017a, 2017b, 2019). Continuing this work, I analyze vowel harmony using phonotactic restrictions over multi-tiered ARs. These restrictions take the form of a constraint which contains a connected piece of a structure—a substructure—that is forbidden, much like the cooccurrence restrictions we often see in phonology. Any structure which contains this forbidden substructure is considered ungrammatical in the same way that a cooccurrence constraint might render a word with a forbidden consonant cluster ungrammatical. The expressivity of such Forbidden Substructure Constraints (FSCs) over multi-tiered ARs encourages us to rethink the complexity of attested vowel harmony patterns.

This chapter explains the formalisms that I use throughout my dissertation. In Section 1 I explain the complexity classes that are relevant for vowel harmony. In section 2 I discuss the precise formulation of the multi-tiered ARs I use. In section 3 I define FSCs over multi-tiered ARs. Section 4 summarizes.

3.1 The Subregular Hierarchy

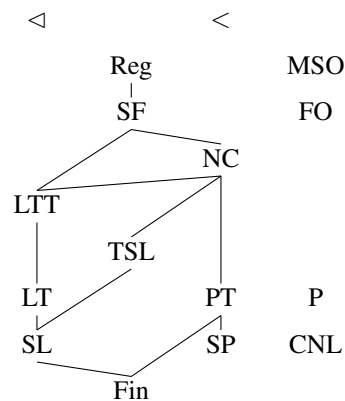
One of the most well-studied data structures is the string. Strings are ordered linear structures generated from a finite set of symbols which is called the alphabet, denoted by Σ . A stringset is a set of strings which makes up a pattern and the set of all possible stringsets is denoted by Σ^* . Different stringsets can be classified according to shared computational properties and these classes are mathematically defined. For example, a significant body of work in computational phonology shows that phonological generalizations are properly contained within the regular class of stringsets (Heinz 2011a, 2011b; Heinz and Idsardi 2013; Kaplan and Kay 1994). Recent work has further established a subregular hierarchy of stringset classes, i.e. star-free (SF) and weaker classes (Heinz, Rawal, and Tanner 2011; Rogers and Pullum 2011; Rogers et al. 2013). The ordering relations relevant for this subregular hierarchy are Successor (\triangleleft) and Precedence (\prec), illustrated for the first element in a string in (29).

(29) Ordering Relations involving a



The successor relation (\triangleleft) refers to strict adjacency, so in (29) the first element is succeeded by only the element one position to its right. While it is not shown, this relation can be extended to ensure that all elements in the string are ordered; so, c would succeed b and d would succeed c. The precedence relation ($<$), on the other hand, refers to general precedence which means that any element in a string precedes all the elements in each position to its right. While (29) shows that a precedes b, c, and d; extending this relation means that b would precede both c and d, and c would precede only d. The subregular hierarchy in (30) classifies stringsets in terms of the relative expressivity of the grammars needed to generate them and is split vertically based on whether the grammars in each class rely on the successor or precedence relation between elements in a string. Each class that is lower on the hierarchy is also a proper subset of the class that dominates it, connected by a straight line.

(30) 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 Piecewise (SP), Finite (Fin)

Ordering Relations: Successor (\triangleleft), Precedence ($<$)

Logical Power: Monadic Second Order (MSO), First Order (FO), Propositional (P), Conjunctions of Negative Literals (CNL)

The diagram in (30) 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 Strictly Local (SL) or Strictly Piecewise (SP) and they generate patterns that can be described using conjunctions of negative literals (CNLs). The FSCs discussed above are a type of negative literal and so the FSCs used to describe a pattern are SL if they are interpreted with the successor ordering relation, as in (31).

- (31) SL_2 CNL
 –CC (commonly known as *CC)

The CNL in (31) illustrates an FSC that forbids two consonants. It is SL_2 because it forbids only two elements and is interpreted with the successor ordering relation. This CNL accurately describes the phonotactic restriction found in some languages which allow only syllables with alternating consonants and vowels (CV or VC) and forbids consonant clusters. As it is written this FSC forbids successive consonants both within and across syllable boundaries.

This FLT approach provides explicit ways of determining the computational complexity of vowel harmony patterns over strings. In this dissertation I adapt this approach to assess the complexity of vowel harmony patterns over multi-tiered ARs. By using FSCs with the successor ordering relation my FLT approach provides the tools necessary to investigate the locality of vowel harmony patterns over both strings and ARs. Lastly, I use this FLT approach to investigate whether the surface well-formedness theory I propose must be restricted to avoid predicting unattested patterns.

Heinz (2010) proposes a “subregular hypothesis”, which claims that all phonological patterns are subregular. Heinz (2018)’s theory of phonotactics further claims that all cooccurrence restrictions can be described using CNLs over strings and are thus SL or SP. However, given evidence of an existing vowel harmony pattern that is necessarily LT over strings in Tutrugbu, I propose a variation of Heinz (2018)’s claim and call it the “ ASL^{VH} hypothesis”: using evidence from vowel harmony patterns found in languages around the world I argue that phonotactic cooccurrence restrictions are describable using a CNL of FSCs over multi-tiered ARs. I further propose a new complexity class called Autosegmental Strictly Local (ASL^{VH}) which includes vowel harmony patterns that are local over ARs but fit into different classes on the subregular hierarchy in (30) when they are represented over strings.

In the following subsections, I will explain further details of the subregular classes outlined here (SL, SP, LT, TSL, and ASL^{VH}) and formally define the differences between CNLs and propositional logic.

3.1.1 Strictly Local

One type of phonotactic constraint 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 (Prince and

Smolensky 1993) markedness constraints with an asterisk (*), and many such markedness constraints can be easily translated into the lowest level of logic in (30). Conjunctions of negative literals (CNLs) are formally defined in (32).

(32) Definition of Conjunctions of Negative Literals (CNLs)

a. Literal: A literal is the most basic kind of propositional statement which represents the presence of a single substructure.

b. CNL: A propositional statement is a CNL if it has the following structure:

$$\neg w_0 \wedge \neg w_1 \wedge \neg w_2 \wedge \dots \wedge \neg w_n$$

where $w_0, w_1, w_2, \dots, w_n$ are literals

In short a CNL is exactly what it sounds like: a set of negated literals that are connected with boolean conjunction. So a set of FSCs is a type of CNL because it consists of a set of literals (substructures) over some representation which are negated (i.e. forbidden) and connected with boolean conjunction. For example, some languages use only syllables that contain one consonant and one vowel (CV or VC) or only a single vowel (V). These languages do not allow syllables to contain clusters of more than one adjacent consonants in either the onset or coda (*CC). The allowed syllable shapes and the restriction on syllable shapes that I just described are summarized in (33). The SL_2 grammar that generates a pattern of syllables which excludes consonant clusters would look like (34).

(33) Syllable shapes

a. Allowed: V, CV, VC

b. Restriction: *CC

(34) SL_2 grammar that forbids

consonant clusters

($\Sigma = \{C, V\}$), CNL: $\neg CC$

The restriction on consonant clusters can be described by the OT style phonotactic constraint in (33b) which forbids two consonants from occurring in succession. This same constraint can also be translated into a single FSC $\neg CC$, which makes a CNL with only one literal. This CNL would read as “Strings which do not contain consonant clusters are well formed”.

SL stringsets are also defined by the property of Suffix Substitution Closure (SSC), which is defined in (35) (Heinz 2018; McNaughton and Papert 1971; Rogers et al. 2013). SSC states that if two strings are in a stringset and both contain the substring x of length $k-1$ —where k is equal to the number of elements in each negative literal (NL)—then swapping the substrings that follow x in each would create two new strings, which are also in the same stringset.

- (35) **Suffix Substitution Closure (SSC):** A stringset \mathbb{L} is strictly k -local iff
 for any string x of length $k-1$ and any strings v, w, y, z ,
 $wxy \in \mathbb{L}$ and $vzx \in \mathbb{L}$ implies that $wxz \in \mathbb{L}$ and $vxy \in \mathbb{L}$

SSC can be used as a test to prove when a stringset is not SL because the definition in (35) states that “a stringset is strictly k -local iff for any string $x \dots$ ”. So a stringset is provably *not* SL_k if it contains a single string for which SSC does not hold. Heinz (2010) provides one example of such a pattern in Navajo, which utilizes consonant harmony between sibilants with respect to the feature [anterior]. So words can contain [s, z] or [ʃ, ʒ], but not both and this restriction applies regardless of how many segments occur between the sibilants. So if Navajo words are represented as strings of segments in the order in which they occur in actual words, the set of Navajo strings is represented as \mathbb{L}_N .

- (36) Navajo (Athabaskan, Na-Dene; Heinz 2010)
- ʃite:ʒ ‘they (dual) are lying’
 - dasdo:lis ‘he (4th) has his foot raised’
 - ʃiyiʃ ‘it is bent, curved’
 - naʃyatʃ ‘I killed them again’

Table 3: SSC test for Navajo


$x = \text{o}^{k-1}$	$wxy = \text{ʃo}^{k-1}\text{ʒ}$	$\in \mathbb{L}_N$
	$vzx = \text{so}^{k-1}\text{z}$	$\in \mathbb{L}_N$
	$wxz = \text{ʃo}^{k-1}\text{z}$	$\notin \mathbb{L}_N$

The test in Table 3 shows that over strings of segments the stringset \mathbb{L}_N representing Navajo sibilant harmony contains two strings $\text{ʃo}^{k-1}\text{ʒ}$ and so^{k-1}z , but when their suffixes—[ʒ] and [z], respectively—are swapped the resulting strings $\text{ʃo}^{k-1}\text{z}$ and $\text{so}^{k-1}\text{ʒ}$ are not in \mathbb{L}_N , which means they are not grammatical strings of Navajo.

3.1.2 Strictly Piecewise

Similar to the SL stringsets, a stringset is Strictly Piecewise (SP) if it is both describable by a CNL interpreted with the precedence ordering relation ($<$) and obeys Subsequence Closure (SC). For example, the Navajo sibilant harmony pattern introduced above and summarized in (36) was proven not to be SL but it is describable with a SP_2 CNL interpreted with precedence in (38).

(37) [ʃite:ʒ] ‘they (dual) are lying’



(38) SP_2 grammar for Navajo harmony

$$\langle \Sigma = \{IPA\}, < \rangle, \text{CNL: } \neg s\text{f} \wedge \neg s\text{ʒ} \wedge \neg z\text{f} \wedge \\ \neg z\text{ʒ} \wedge \neg \text{f}s \wedge \neg \text{f}z \wedge \neg \text{ʒ}s \wedge \neg \text{ʒ}z$$

Given an alphabet (Σ) of IPA symbols, the CNL for Navajo sibilant harmony in (38) is SP_2 over strings of segments because each FSC contains only two sibilants. The string in (37) shows that each segment is succeeded by a single segment to the right, but the SP_2 FSCs refer to sibilants like the first [ʃ] and last [ʒ], which are connected by a precedence relation; so the string in (37) is grammatical regardless of how many segments intervene between the sibilants. In other words, the FSCs forbid the segment sequences listed in (38) from occurring within a grammatical Navajo string regardless of how many segments might occur between them.

SP stringsets are also defined by the property of Subsequence Closure (SC) defined in (39) (Heinz 2018; Rogers et al. 2010, 2013). SC states that if a string w is in a stringset and v is a subsequence of or equal to w ($v \sqsubseteq w$), then v is also in the stringset.

(39) **Subsequence Closure (SC):** A stringset \mathbb{L} is strictly k -piecewise iff

for every string $w \in \mathbb{L}$, if $v \sqsubseteq w$ then $v \in \mathbb{L}$

Similar to SSC, the definition of SC in (39) states that “a stringset... is strictly k -piecewise iff for every string...”. This means that if a single string can be found within a stringset for which SC does not hold then that stringset is provably *not* SP_k . In other words SC can be used as a test to prove when a stringset is not SP_k . One example of such a pattern is called First-Last Harmony (FLH) and has been claimed to be unattested. A FLH pattern requires that the first element in a string and the last element in a string share some property regardless of what comes between them. Lai (2015) provides an example of a FLH pattern based on Navajo sibilant harmony, discussed above. In Navajo all sibilants must have the same anteriority regardless of where they occur in a word and how many segments occur between them. In a FLH pattern, only the first and last sibilant would have to have the same anteriority regardless of what segments (even sibilants) might occur between them. The set of strings that make up such a FLH pattern is represented as \mathbb{L}_{FLH} .

(40) Hypothetical FLH data

- a. sotofotos
- b. ʃisizeʒ

Table 4: SC test for First-Last harmony (FLH)

$$\frac{w = \text{sotofotos} \in \mathbb{L}_{FLH}}{v = \text{fs} \sqsubseteq w \notin \mathbb{L}_{FLH}}$$

The test in Table 4 shows that over strings of segments the stringset \mathbb{L}_{FLH} representing the sibilant FLH pattern contains a grammatical string which itself contains a substring that is not in \mathbb{L}_{FLH} and is thus not a string with First-Last harmony.

3.1.3 Locally Testable

Unlike the phonotactic restrictions modeled in (34) and (38), Heinz (2018) argues that certain logically possible phonological patterns are unattested because they must be categorized in a higher subregular stringset class and their description requires a more expressive logic. Statements made using CNLs as they are defined above in (32) form a strict subclass of the possible statements made using full propositional logic, defined in (41). CNLs use the boolean connectives negation and conjunction but only use them to relate literals, which are defined in (41a) and each represent the presence of a single substructure. On the other hand, full propositional logic is more expressive than CNLs because it includes the full set of boolean connectives (negation, conjunction, disjunction, and conditional) which can be used to relate either literals or entire statements. Propositional logic is also made up of statements that can be true or false and is denoted as \mathcal{L}^P . The propositional logic defined in (41) includes a syntax which defines the set of possible statements; the semantics is defined as usual.

(41) Definition of propositional logic

Syntax: Given a statement ϕ , $\phi \in \mathcal{L}^P$ if:

- a. $\phi = u$ for a substring u of some word $w \in \times \Sigma^* \times$ (literal)
- b. $\phi = (\neg \psi)$ for some $\psi \in \mathcal{L}^P$ (negation)
- c. $\phi = (\psi_1 \wedge \psi_2)$ for some $\psi_1, \psi_2 \in \mathcal{L}^P$ (conjunction)
- d. $\phi = (\psi_1 \vee \psi_2)$ for some $\psi_1, \psi_2 \in \mathcal{L}^P$ (disjunction)
- e. $\phi = (\psi_1 \Rightarrow \psi_2)$ for some $\psi_1, \psi_2 \in \mathcal{L}^P$ (conditional)
- f. and nothing else is in \mathcal{L}^P

The propositional logic defined in (41) is expressive enough to describe the FLH pattern introduced above. To reiterate, according to Heinz (2018) a FLH pattern is one in which the first and last element of any string must be the same regardless of what occurs in between them and (42) utilizes the diagnostics described above to prove that a FLH pattern over strings (\mathbb{L}_{FLH}) is neither SL nor SP.

(42) First-Last Harmony (FLH)

- a. not SL: $so^{k-1}s \in \mathbb{L}_{FLH}$ and $fo^{k-1}f \in \mathbb{L}_{FLH}$, but $so^{k-1}f \notin \mathbb{L}_{FLH}$ and $fo^{k-1}s \notin \mathbb{L}_{FLH}$
- b. not SP: $sotofotos \in \mathbb{L}_{FLH}$ and $fotos \sqsubseteq sotofotos$, but $fotos \notin \mathbb{L}_{FLH}$

In (42a) application of SSC clearly illustrates that \mathbb{L}_{FLH} cannot be SL_k for any k since neither of the strings that result from substituting suffixes would be included in the \mathbb{L}_{FLH} stringset. In (42b) application of SC clearly shows that \mathbb{L}_{FLH} cannot be SP since a string in \mathbb{L}_{FLH} can also contain a subsequence that is not a possible string in \mathbb{L}_{FLH} . This pattern can be described using a conjunction of conditionals from the propositional logic defined in (41).

(43) Propositional logic statement of \mathbb{L}_{FLH} :

$$\begin{aligned} & (\#s \Rightarrow \neg\#) \wedge (\#s \Rightarrow \neg\mathfrak{z}\#) \wedge (\#z \Rightarrow \neg\mathfrak{z}\#) \wedge (\#z \Rightarrow \neg\#) \wedge \\ & (\#f \Rightarrow \neg\#) \wedge (\#f \Rightarrow \neg\mathfrak{z}\#) \wedge (\#\mathfrak{z} \Rightarrow \neg\mathfrak{z}\#) \wedge (\#\mathfrak{z} \Rightarrow \neg\#) \end{aligned}$$

The conjunction of conditionals in (43) describes the following pattern: only strings in which the string-initial and string-final sibilant have the same anteriority are grammatical in \mathbb{L}_{FLH} . The string edge boundaries are represented by the # symbol.

Rogers et al. (2013) claims another defining property of LT stringsets is Local Test Invariance (LTI), which states that two strings with the same set of k -factors are either both members of \mathbb{L} or neither is. According to Rogers et al. (2013) “. . . the substrings that occur within a string are referred to as its factors” and so a k -factor is a substring of length k .

(44) **Local Test Invariance (LTI)**: A stringset is LT iff there is some k such that, for all strings x and y , if $\times \cdot x \cdot \times$ and $\times \cdot y \cdot \times$ have exactly the same set of k -factors then either both x and y are members of the stringset or neither is.

Like SSC and SC, LTI can be used as a test to prove when a pattern is not LT. The FLH pattern described in this section is properly LT because it is describable with a propositional logic statement and is provably not SL or SP; FLH fails the SSC and SC tests but not LTI.

3.1.4 Tier-based Strictly Local

Heinz, Rawal, and Tanner (2011) argue for a tier-based class of stringsets to account for long-distance patterns like vowel harmony; Aks nova (2017) and Aks nova and Deshmukh (2018) agree that long-distance vowel harmony patterns are TSL. A TSL grammar is a phonotactic grammar in which a subset of elements in a string are projected onto a separate tier; in other words a tier projection can only contain symbols from the

original string’s alphabet. Different TSL tiers can have disjoint sets of elements but no TSL 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, i.e. vowels. 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” (Aksënova 2017). The examples below are taken from Akinlabi (2009).

(45) Lokaa (Niger-Congo; Akinlabi 2009)

- a. èsìsòn ‘smoke’, *èsìsòn
- b. èsìsòn ‘housefly’, *èsìsòn
- c. lèjìmè ‘matriclan’, *lèjìmà
- d. ékílikà ‘kind of plant’, *ékílikà

(46) ATR harmony in Lokaa



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 (46). 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. Thus over the non-high vowel tier projection Lokaa ATR harmony also obeys SSC. The TSL grammar that generates the Lokaa vowel harmony pattern would look like (47).

(47) TSL grammar for Lokaa harmony

- a. Tier of non-high vowels: $T = \varepsilon, e, o, \text{ə}, \text{ɔ}, a$
- b. CNL: $\neg[\alpha\text{ATR}][\beta\text{ATR}] = \neg\varepsilon\varepsilon \wedge \neg e\varepsilon \wedge \neg\varepsilon o \wedge \neg o\varepsilon \wedge \neg e\text{ə} \wedge \neg \text{ə}e \wedge \neg e\text{ɔ} \wedge \neg \text{ɔ}e \wedge \neg e\text{ə} \wedge \neg \text{ə}o \wedge \neg o\text{ə} \wedge \neg \text{ə}o \wedge \neg o\text{ɔ} \wedge \neg \text{ɔ}o \wedge \neg \text{ə}a \wedge \neg a\text{ə} \wedge \neg a\text{o} \wedge \neg o\text{a} \wedge \neg a\varepsilon \wedge \neg \varepsilon a$

The TSL grammar in (47) illustrates how 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 new tier projection thus puts harmonizing [-high] vowels in succession when they were separated by one or more transparent vowels on the original string. The harmony pattern can now be described locally by FSCs which forbid successive [-high] vowels with different ATR feature values.

3.1.5 Autosegmental Strictly Local

Rogers et al. (2013) discussed how the computational complexity of a pattern depends upon the set of representational primitives over which a pattern is evaluated. For example, a phonological pattern can provably fall into one of the complexity classes discussed above when it is represented over strings. However, utilizing a more complex representation reduces the computational power and logical expressivity needed to evaluate a pattern. Jardine (2019) thus proposed a new complexity class for tone patterns in which a pattern is considered local when it is represented over ARs, called Autosegmental Strictly Local (ASL^T). ASL^T ARsets are describable by a CNL of connected FSCs over ARs, but no test has yet been developed to mathematically prove when a set of ARs is not ASL^T . Much previous work also used ARs to represent vowel harmony patterns (Clements 1976; Goldsmith 1976; McCarthy 1988; Padgett 2002; Sagey 1986; van der Hulst 2017; Walker 2010, 2011, 2014b). I follow Jardine (2016b) by comparing autosegmental vowel harmony patterns to their analogous stringset patterns. Using the mathematical tests described above I determine that when evaluated over strings vowel harmony patterns span multiple complexity classes such as SL, SP, TSL, and LT. I extend Jardine (2019)'s ASL class to outline an ASL^{VH} class of ARsets which encompasses vowel harmony patterns that fall into different complexity classes over strings but are all describable with a CNL of FSCs over multi-tiered ARs. Thus evidence from vowel harmony supports my ASL^{VH} hypothesis.

This dissertation includes new findings which provide evidence for the expressivity and restrictiveness of the new ASL^{VH} class. First, Heinz (2018) and Lai (2015) claimed that first-last harmony (FLH) patterns are unattested but Walker (2011) analyzes Eastern Meadow Mari (EMM) which she described as utilizing a FLH pattern. While I prove in chapter 6 that EMM is not strictly a FLH pattern it is nonlocal for a similar reason; both are patterns which cannot be described locally and thus do not fit into the ASL^{VH} class. Jardine (2019) defines ASL patterns as being forbidden substructure grammars (FSG) over ARs. First he proves that for each FSG there is an equivalent statement in first order logic (FO). FO is the level of logic which describes exactly the Star Free (SF) stringsets. Jardine (2019) then goes on to prove that the ASL class is a strict subset of the SF class because all ASL patterns are also describable as SF patterns but not all SF patterns are describable as ASL patterns.

Some SF patterns cannot be ASL because the constraints needed to describe them would not form connected substructures. For example given the context of back harmony, in order to describe a FLH pattern like the one described in (40) and (42)-(43) an ASL grammar requires the existence of some k such that there is a set B of forbidden substructures of length k which ban all and only ARs of the forms $\#[-back]w[+back]\#$ or $\#[+back]w[-back]\#$ for some $w \in \Sigma$. This means banning exactly the set of ARs which contain one of the structures in (48). Both w and the ‘...’ in (48) represent the string of features which intervene between the

first and last feature on the back tier in this example.


- (48) a. $\# \rightarrow -\text{back} \rightarrow \dots \rightarrow +\text{back} \rightarrow \#$ b. $\# \rightarrow +\text{back} \rightarrow \dots \rightarrow -\text{back} \rightarrow \#$

Because B must contain connected substructures there is no finite set B and no finite k which can ban all such ARs. Therefore, the FLH pattern is SF and cannot be ASL^{VH} . While all ASL patterns are also describable as SF patterns the existence of a SF pattern like FLH that is not describable as an ASL pattern proves that $\text{ASL}^{VH} \subsetneq \text{SF}$. In addition, the attested positional transparency pattern in EMM is properly SF and the unbounded circumambient (UCA) blocking pattern in Tutrugbu is properly LT over strings. The existence of a properly SF pattern and a properly LT pattern amongst the set of attested vowel harmony patterns subverts Heinz (2018)'s claim that only SL or SP phonotactic patterns are attested.

Second, when transparent vowels do not contrast in the harmonic feature, multi-tiered ARs represent the vowel harmony pattern in a local way without relying on underspecification. When looking at a surface string, transparency makes it look like harmony has skipped over certain vowels and the number of transparent vowels determines the distance between harmonizing vowels. However, with fully specified ARs all successive transparent vowels are associated to a single feature. This means that no matter how many successive transparent vowels are present in a string, the harmonizing feature values are only separated by a single feature on the same tier, as shown in (49) below.

- (49) Local transparency in Finnish:

[*maisemia*] ‘scenery.plural.partitive’ (Sulkala and Karjalainen, 1992)




The multiple association in (49) makes transparency local over fully specified ARs. In addition, the fact that transparent vowels without a harmonic contrast and a LT harmony pattern both become local when evaluated over ARs provides strong evidence for my ASL^{VH} hypothesis and the expressivity of the ASL^{VH} class. The exclusion of an unattested properly SF string pattern like FLH provides strong evidence for the restrictiveness of the ASL^{VH} class. The exclusion of an attested SF pattern like EMM suggests that there is still room for further investigation into the expressivity of ARs.

3.1.6 Types of tiers

When comparing the analyses of vowel harmony patterns over strings with ARs one might want to ask the following question: What is the difference between a TSL and an autosegmental tier? A TSL 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 a TSL tier is directly associated to the identical segment in the original string and no other. TSL tiers cannot utilize multiple association. On the other hand, autosegmental tiers consist of a subsegmental feature, such as backness, height, ATR, or tone. For vowel harmony patterns each vowel in a word is associated to an element on multiple feature tiers because a vowel is characterized by multiple different features and each feature is represented on a separate tier. In addition, each feature on a given autosegmental tier can be associated to multiple vowels, as shown in the Lokaa example in (50).

(50) ATR harmony in Lokaa



Unlike TSL tiers, autosegmental tiers consist of elements in a disjoint set from the set of elements to which they are associated and they can be multiply associated. Whether it be tones multiply associated to tone bearing units (TBUs) or vowel features multiply associated to vowels, the possibility for multiple association crucially distinguishes autosegmental tiers from TSL tiers.

By using a more complex representation—one with autosegmental tiers—the complexity of vowel harmony patterns is reduced. By comparing vowel harmony patterns over ARs with their analogous patterns over strings I am able to determine where vowel harmony patterns fit within the subregular computational complexity hierarchy of (30). I show that vowel harmony patterns of various complexities over strings can be captured using FSCs over ARs with multiple tiers. With this strong evidence I thus propose the new ASL^{VH} class which encompasses all such vowel harmony patterns and supports the ASL^{VH} hypothesis that all phonotactic patterns are describable with a CNL of FSCs over ARs.


3.2 Autosegmental Representations (ARs)

This section outlines the basic assumptions and definitions needed to understand how multi-tiered ARs are used to represent vowel harmony patterns. Multi-tiered ARs decrease the complexity of vowel harmony patterns. Enriching the representation to use ARs allows a broad variety of patterns to be grouped together within the

same class because they can be described locally. I determine that multi-tiered ARs adequately capture vowel harmony patterns and so their expressivity can be compared to ARs for other types of phonological patterns, such as tone.

3.2.1 Multi-tiered ARs

In the last decade, a body of work has begun to emerge which takes a formal perspective on an investigation of the range of patterns that can be represented using autosegmental tiers. I build on the work begun by Jardine (2016b), Jardine (2017b), and Jardine (2019) to investigate the expressive power needed to represent one such set of patterns called vowel harmony. Vowel harmony patterns generally refer to subsegmental features, such as backness, height, or ATR. I adopt a “bottlebrush” feature representation in which each feature—with a + or - value—is represented on a separate tier and is associated directly to a vowel (Clements 1976; Hayes 1990; McCarthy 1988; Padgett 1995); thus each vowel is associated to multiple featural tiers. Such ARs include at least one additional tier compared with ARs of tone patterns, which utilize only two tiers (Jardine 2016b, 2017b, 2019). For example, assuming binary features, vowel features like $[\pm \text{back}]$, $[\pm \text{high}]$, etc. are represented on separate tiers and each is associated to a vowel.



The examples in (51) clearly illustrate that association relations are represented with straight lines that connect elements on different tiers, i.e. vowels and features. Example (51b) illustrates that the successor ordering relation between elements on a tier is represented with an arrow.

While there are phonological patterns in which vowels and consonants can interact (Padgett 2011), vowel harmony patterns do not generally include such interactions and are able to skip consonants and operate over vowels only. Some theories of feature geometry have claimed that this is possible because vowel features and consonant features are distinguished by including a vocalic node, labeled as “Vocalic” or “V-place”, to which all vocalic features are associated (Clements 1991; NiChiosain and Padgett 1993; Padgett 1995, 2002). Baković (2000) effectively “ignore[d] consonants altogether” and assumed that vowels “are adjacent across intervening consonants”. In order to explain this somewhat long-distance behavior of vowel harmony Nevins (2010) introduced a notion of “relativized locality” that is similar to the TSL notion of locality on a tier. In this dissertation I formalize these observations about vowel harmony by abstracting away from vowel-consonant interactions and assuming that consonants cannot be associated to vowel features or vice versa. The ARs I


use effectively combine a TSL notion of a vowel tier projection with a bottlebrush feature representation to illustrate that only vowel features are relevant for vowel harmony. In short, a string of vowels is represented on a single tier with the successor relation between vowels; and vowel features are associated directly to elements on the vowel tier. Elements on the vowel tier are represented as V except where IPA symbols provide a visual aid to connect an AR to a specific data point. Using IPA also avoids the visual clutter of a full bottlebrush representation with all the vowel feature tiers by allowing ARs to show only the features relevant for the pattern and discussion at hand. Including a tier with only a string of vowels in these multi-tiered ARs further aids the visual comparison of AR analyses to string-based analyses of vowel harmony patterns.

3.2.2 Representational assumptions

Use of ARs requires discussion 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 (52)-(55) below.

First, FS means that each featural element must be associated to at least one vowel and each vowel 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.


(52) Violates FS



The hypothetical representation in (52) straightforwardly violates FS because the second vowel is not associated to any feature on the ATR tier.


Second, the NCC states that association lines between the vowel tier and a feature tier never cross. Odden (1994) adds that the NCC can only evaluate an association between the vowel and one featural tier at a time. The representation in (53) violates the NCC because +ATR precedes -ATR, but is associated to a vowel that succeeds a vowel associated to -ATR; this configuration creates visually crossed association lines.

(53) Violates NCC



A notable effect of FS along with the NCC is that they prevent what have been called gapped structures (Archangeli and Pulleyblank 1994; Ringen and 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.


(54) Gapped Structures



As shown in (54a), FS would prevent gapped structures in which the “skipped” vowel is not associated to anything on that particular feature’s tier. In (54b), 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 (55) violates the OCP because on both the ATR and low feature tiers there are two identical successive features, -ATR and -low respectively.

(55) Violates OCP




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 (56).

Both the NCC and the OCP have also been derived via a concatenation operation (\circ) that merges autosegmental “graph primitives” (Jardine and Heinz 2015). 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 on the right of (56) is derived from the set of graph primitives on the left. Each primitive 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 shown in (56). The merging of identical adjacent features essentially prevents surface ARs from having multiple successive features with the same value and crossed associations, thus satisfying both the OCP and the NCC. However, if two vocalic 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.

(56) Concatenation of adjacent autosegmental graph primitives satisfies FS, NCC, and OCP



The AR in (56) 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. This AR 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.

3.3 Definition of Constraints

In this dissertation I use forbidden substructure constraints (FSCs) to analyze vowel harmony as generated by a set of phonotactic restrictions (Nevins 2010) and to determine the locality of such restrictions over multi-tiered ARs. Previous work on the logical descriptions of formal languages and their applications to phonological well-formedness constraints (Heinz, Rawal, and Tanner 2011; Rogers et al. 2013) led to the development of the theory of a forbidden substructure grammar (FSG) (Jardine 2017b). A FSG is a CNL of the form in (57) below; such a grammar will generate a set of well-formed structures that does not contain any of r_1 through r_n .

(57) Forbidden substructure grammar (Jardine, 2017)

$$\neg r_1 \wedge \neg r_2 \wedge \neg r_3 \wedge \dots \wedge \neg r_n$$

Each r in (57) represents a single literal—e.g. a connected substructure—that is forbidden. If a full structure—i.e. a string or AR—contains one of these substructures (r_1 – r_n) the entire structure is ungrammatical. FSCs can be written using an asterisk (*) to denote that a piece of structure is marked or a negation symbol (–) as is used in logical statements like the FSG in (57).

FSCs serve as a type of phonotactic restriction such that “well-formedness is based on contiguous structures of a specific size” (Jardine 2017b). 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 FSG because it can be captured with FSCs by referring to a subset of the elements within structures and their connections. Jardine (2017b) uses FSCs to show that attested tone patterns are local in this way. In this dissertation I utilize FSCs over multi-tiered ARs to argue that a variety of vowel harmony patterns are local in the same way.

Over multi-tiered ARs, the forbidden piece of structure necessarily includes elements on more than one tier and their connections. Because FSCs over multi-tiered ARs can include an element on the vocalic tier as well as its connections to elements on other tiers vowel harmony patterns can utilize a feature on one tier in order to constrain the harmony that occurs on another tier. For example, in Turkish round harmony is said to be parasitic on height because only [+high] suffix vowels assimilate in roundness to the closest vowel to its left.

The vowel inventory of Turkish in Table 5 consists of eight vowels with three main featural distinctions: \pm high, \pm back, \pm round. The [+high, -round] vowels are [i, i], the [+high, +round] vowels are [y, u], the [-high, -round] vowels are [e, a], and the [-high, +round] vowels are [ø, o].

Table 5: Turkish Vowels

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

The surface generalization for rounding harmony in Turkish can be described as follows: a [+high] suffix vowel has the same round feature as the closest vowel to its left, but a [-high] suffix vowel does not. The data in Table 6 below is adapted from Padgett (2002) and illustrates this Turkish round harmony generalization. The first four nominative singular roots contain [-round] vowels while the last four contain [+round] vowels. The genitive suffix contains a [+high] vowel and so when it follows one of the first four roots it surfaces as [-round] and when it follows one of the last four roots it surfaces as [+round]. The nominative plural suffix, on the other hand, contains a [-high] vowel which remains [-round] regardless of which root precedes it.

Table 6: Turkish round harmony data from Padgett (2002)


	a. Nom.sg	b. Gen.sg	c. Nom.pl	Gloss
1.	ip	ip+in	ip+ler	'rope'
2.	kiz	kiz+in	kiz+lar	'girl'
3.	el	el+in	el+ler	'hand'
4.	sap	sap+in	sap+lar	'stalk'
5.	jyz	jyz+yn	jyz+ler (*lør)	'face'
6.	pul	pul+un	pul+lar (*lør)	'stamp'
7.	køj	køj+yn	køj+ler (*lør)	'village'
8.	son	son+un	son+lar (*lør)	'end'

Some Turkish roots contain more than one syllable and do not necessarily obey rounding harmony (Nevins 2010). These disharmonic roots contrast with the disyllabic words in Table 6 because they include two [+high] vowels, which are each associated to different features on the round tier. This particular vowel combination in the AR below is specifically what the surface generalization of Turkish round harmony should forbid across a morpheme boundary, but within roots Turkish allows it.

(58) Turkish disharmonic roots (Nevins 2010)

- a. butik 'boutique'
- b. kuvvet 'strength'


(59) [butik] 'boutique'



Since Turkish includes disharmonic roots like (59), roots and suffixes must be differentiated within ARs. However, the following analysis is based solely on the data provided in Table 6 and does not take into account additional possible Turkish morphemes. I thus distinguish between root and suffix morphemes based on their relative position with respect to a morpheme boundary primitive: root-final vowels are succeeded by a morpheme boundary and a suffix succeeds the morpheme boundary. Morpheme boundaries are included as an additional element on the vocalic tier and are represented with '+'.

The surface generalization of the Turkish round harmony pattern can be described using the three FSCs in (60). Each of these FSCs constitutes a single autosegmental substructure which is forbidden from occurring within a full AR just like r_1-r_n in (57).

(60) Turkish round harmony FSCs



The conjunction of these three FSCs thus constitutes a CNL of the form in (57) and so (60) is a FSG which describes Turkish round harmony. The first two FSCs in (60a)-(60b) forbid a [+high] suffix vowel from being associated to a round feature with a different value from the round feature associated to the root-final vowel. In other words, the first two FSCs enforce rounding harmony with [+high] suffixes. The third FSC in (60c) forbids a [-high] suffix vowel from being associated to the same [+round] feature as the root-final vowel. This last FSC describes the fact that on the surface [-high] suffix vowels are not [+round] even when the root-final vowel is.

The Turkish genitive singular suffix shown in Table 6 clearly illustrates the effect of round harmony on [+high] suffix vowels. For example, the nominative singular root [jyz] contains a [+high, +round] vowel and so the [+high] genitive singular suffix [jyz+yn] is also [+round]. The surface ARs for the grammatical [jyz+yn] ‘face (Gen.sg)’ and an ungrammatical variant of it *[jyz+in] are shown in (61) below.

(61) Turkish rounding harmony with [+high] suffix

a. [jyz+yn] ‘face (Gen.sg)’



b. *[jyz+in]



The AR in (61a) illustrates that a [+high] suffix vowel which is associated to the same round feature as the root-final vowel is grammatical because it does not contain the forbidden substructure of (60a). On the other hand, the AR in (61b) is ungrammatical because the [+high] suffix vowel is associated to a round feature with a different value than the round feature associated to the root-final vowel and thus it contains the forbidden substructure of (60a).


In addition, the Turkish nominative singular suffix shown in Table 6 provides a clear example of how [-high] suffix vowels do not harmonize with the root-final vowel with respect to rounding. The nominative singular root 7a in Table 6 contains a [-high, +round] vowel and so the [-high] nominative plural suffix vowel in 7c is [-round]. The surface ARs for 7c [køj+ler] ‘village (Nom.pl)’ along with its ungrammatical counterpart *[køj+lor] are shown in (62) below.

(62) Turkish rounding harmony with [-high] suffix

a. [køj+ler] ‘village (Nom.pl)’




b. *[køj+lør]



The AR in (62a) shows that a grammatical word with a [-high] suffix vowel does not include rounding harmony like words with [+high] suffix vowels do. This AR is grammatical because it does not contain the forbidden substructure of (60c). By contrast, an AR with a [-high] suffix vowel that is associated to the same [+round] feature as the root-final vowel like (62b) is ungrammatical because it contains the forbidden substructure of (60c).

While a [-high] suffix vowel cannot be associated the same [+round] feature as the root-final vowel, it can be associated to the same [-round] feature as a root-final vowel. This possibility is exemplified by the grammatical Turkish word in Table 6 example 1c and the AR is shown below in (63).

(63) [ip+ler] ‘rope (nom.pl)’



This AR is grammatical because the FSG for Turkish round harmony includes (60c) and does not include a constraint that forbids a [-high] suffix vowel from being associated to the same [-round] feature as the root-final vowel. In other words, a Turkish suffix vowel can only be associated to [+round] if it is also associated to [+high]. According to the generalization presented in Table 6 the [-high] suffix vowels can only be associated to [-round] but never [+round].

Turkish rounding harmony shows that FSCs over multi-tiered ARs are a useful form of phonotactic constraint because the FSG that describes Turkish rounding harmony consists of a set of FSCs. The Turkish FSCs in (60) are negative constraints because they describe pieces of autosegmental structure that are forbidden. The conjunction of these FSCs over multi-tiered ARs thus forms a CNL of the same form as (57) which can generate a vowel harmony pattern that is Autosegmental Strictly Local (ASL^{VH}).

3.4 Summary

In this chapter I defined the formalisms used throughout my dissertation. As a theoretical framework, FLT provides precise methods for evaluating the computational complexity of surface patterns; I extended previous work to situate vowel harmony within the complexity hierarchy that has been proposed for phonology. I precisely defined the multi-tiered ARs whose expressivity I argue allows vowel harmony patterns to be considered local when analyzed as surface phonotactic restrictions. More specifically, autosegmental tiers allow patterns to be represented locally in a way that TSL tiers cannot. I thus propose a new complexity class

called ASL^{VH} which cross-cuts the string-based subregular hierarchy because it contains vowel harmony patterns from a variety of stringset classes that are all local over multi-tiered ARs and excludes the unattested First-Last Harmony (FLH) pattern. The remainder of this dissertation investigates the breadth of this new ASL^{VH} class. I analyze a variety of vowel harmony patterns and show that they are local over multi-tiered ARs. The representations defined in this chapter provide evidence for my ASL hypothesis: vowel harmony patterns are describable using a CNL of FSCs over multi-tiered ARs and are thus ASL^{VH} .

4 Neutral vowels are simple


This chapter begins my investigation of the locality of vowel harmony with neutral vowels. In some well known vowel harmony patterns there are vowels that do not perpetuate or do not undergo vowel harmony, but nonetheless they hold a special status; such vowels can either block or be transparent to the assimilation of a vowel feature. Blocking vowels may appear to have assimilated on the surface and these have been called icy targets (Jurgec 2011); or they can have a different harmonic feature value and these have been called opaque. On the surface, regardless of the blocker's harmonic feature value it crucially requires the vowels on either side of it to have different values for the harmonic feature. Transparent vowels have the opposite effect in that they have not assimilated but vowels on either side of a span of transparent vowels must have the same harmonic feature value on the surface. When grouped together, transparent and blocking vowels are referred to as neutral vowels (Akinlabi 2009; Krämer 2001; Pulleyblank 1996; van der Hulst 2016; van der Hulst and Smith 1986). In this chapter I will show that two well-studied and often-cited vowel harmony patterns with neutral vowels in Akan and Finnish fit into different stringset classes— SL_2 and SP_2 respectively— but both are local over multi-tiered ARs and thus fit into the ASL^{VH} class of ARsets.

The traditional definition of neutral vowels states that they do not have a harmonizing pair (van der Hulst 2016, 2018; van der Hulst and Smith 1986; van der Hulst and van de Weijer 1995) and they can be either transparent or blocking. For example in a backness harmony pattern the neutral vowels might be [-back, -low, -round] and have no [+back, -low, -round] counterpart so whether they are transparent or blocking they could not participate in harmony. However, van der Hulst and Smith (1986) adds to this definition one other type of neutral vowel which “[does] have a harmonic counterpart and hence [does] not have a predictable value for the harmonic feature. . . [but] may behave as transparent or [blocking].” They further specify that “vowels of this type. . . are neutral without there being a neutralization of an opposition” (van der Hulst and Smith 1986, 234). Following van der Hulst and Smith (1986)'s addition to the definition of neutral vowels I discuss patterns in this dissertation which involve vowels that behave as either transparent or blocking and are thus considered neutral. Whether or not there is a neutralization in the harmonic feature for these vowels they exemplify harmony patterns that have always been attributed to neutral vowels: they can either require harmony (transparent) or allow disharmony (blocking) on either side of them.


In this chapter I also explain why blocking and transparent vowels are important for evaluating different representations of vowel harmony assimilation. The literature refers to vowel harmony as a process with two possible assimilation mechanisms: spreading or agreement, but in this dissertation I analyze only surface representations as either diffuse or iterated, shown in (64).

(64) Surface ARs of assimilation mechanisms

a. Diffuse



b. Iterated




A diffuse AR is one in which multiple harmonizing vowels are associated to a single feature, i.e. +F in (64a). An iterated AR is one in which harmonizing vowels are associated to separate iterations of a feature with the same value, as in (64b). A single representation can include both structures, as (64b) does; but a harmony pattern represented by (64b) is considered to be iterative because the harmonizing vowels are associated to separate iterations of one feature value (+F) while the transparent vowels are associated to the opposing feature value (-F) which is diffused between them. So F harmony is iterative when +F is iterated as in (64b) and diffuse when +F is diffused as in (64a).


Vowel harmony also often involves an interaction between more than one feature. While vowels within a word will harmonize in a single feature like F that harmony can also depend on the presence or absence of an association with a separate feature G that is represented on a different tier. Multi-tiered ARs thus also can be diffuse or iterated with respect to the harmonizing feature, as shown in (65).

(65) Multi-tiered Surface ARs of assimilation mechanisms

a. Diffuse



b. Iterated



The feature G in (65b) could also be diffused, but because the harmonizing vowels have assimilated to the F feature the F harmony pattern would still be considered iterative. ATR harmony with blocking vowels in Akan exemplifies vowel harmony represented with diffuse ARs—like in (65a)—and can be captured by a single FSC over a multi-tiered AR. Back harmony with transparent vowels in Finnish exemplifies vowel harmony represented with iterated ARs—like in (65b)—and can be captured by a set of four FSCs over multi-tiered ARs.

Lastly, the presence or absence of a harmonic contrast can have a major effect on the complexity of vowel harmony patterns especially those with transparency. In this chapter I show that the back harmony pattern in Finnish is local over multi-tiered ARs because the transparent vowels do not contrast in backness. In the next chapter I show that the back harmony pattern in Eastern Meadow Mari (EMM) has transparent vowels

which do contrast in backness so it is not local over multi-tiered ARs. Crucially, the locality of these iterative harmony patterns depends upon whether or not a single feature can be diffused across a span of transparent vowels. The non-dominant feature [-back] is diffused when transparent vowels do not contrast as in Finnish and it can be iterated when transparent vowels do contrast as in EMM. The presence or absence of a contrast in backness amongst the transparent vowels thus determines which type of association is possible and whether or not each of these patterns fits into the ASL^{VH} class of ARsets proposed in this dissertation.

In this chapter I show that two vowel harmony patterns with neutral vowels and different assimilation mechanisms are local over multi-tiered ARs because they can be described using FSCs. I further demonstrate that over multi-tiered ARs the ASL^{VH} class includes vowel harmony patterns which fall into different stringset classes—as Jardine (2016b) claimed for two-tiered ARs of tone patterns. In 4.1, I present diffuse ATR harmony with a blocking low vowel in Akan. In 4.2, I present iterated back harmony with transparent vowels in Finnish.

4.1 Blocking vowels: Akan

In this section I discuss an example of blocking in Akan ATR harmony that is often cited in the vowel harmony literature. I begin by translating whole words into strings of vowels and propose a string-based analysis of Akan in order to determine that the Akan pattern is SL_2 over strings of vowels. My string analysis methodology also clarifies that Akan utilizes an ATR contrast even among [+low] blocking vowels and so these blockers are assimilated to the ATR feature of the succeeding vowel. Lastly I present my novel FSC analysis over multi-tiered ARs which shows how these expanded representations describe the Akan pattern in a local way and show that it fits into my new ASL^{VH} class.

One example of vowels that block harmony is found with ATR harmony in Akan (Casali 2012; Clements 1976; Dolphyne 1988; O’Keefe 2004). The Akan surface vowel inventory, in Table 7, consists of ten vowels with two main featural distinctions: \pm ATR and \pm low. There are two [+low] vowels, [æ] and [a], [+ATR] and [-ATR], respectively. Different sources use different symbols to represent the [+ATR, +low] vowel in Akan; Clements (1976) used [ɜ] but the more recent sources O’Keefe (2004) and Casali (2012) use [æ]. According to the 2015 IPA chart [æ] is a mid central vowel so I follow O’Keefe (2004) and Casali (2012) by using [æ] to represent the [+ATR, +low] vowel and [a] to represent the [-ATR, +low] vowel. 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]. The vowel chart in Table 7 contains exactly the two sets of vowels described by O’Keefe (2004). The inclusion of two [+low] blockers—the [+ATR] [æ] and the [-ATR] [a]—may appear to contradict the definition of neutral vowels from the literature which states that they do not have a harmonizing pair (van der Hulst 2016, 2018; van der Hulst and Smith 1986; van der Hulst and van de Weijer 1995), but as I explained

above they are still considered neutral because they are blocking. I thus treat both vowels as equal blockers of ATR harmony in Akan because they are both [+low].

Table 7: Akan Vowels

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

These two [+low] vowels are each succeeded by a vowel associated to the same ATR feature which makes it look like they have both undergone and blocked ATR harmony from their right.

The basic surface ATR harmony generalization in Akan is that if a word contains a sequence of [-low] vowels, then those vowels will also share the same ATR feature. On the surface full ATR harmony thus includes words like those in (66) which contain only [-low] vowels that are all either [+ATR] or [-ATR].⁴

(66) Akan words with only [-low] vowels (Clements 1976)

- a. tie ‘listen’
- b. obejii ‘he came and removed it’
- c. wubenom? ‘you will suck it’
- d. ɔbɛjɛɪ ‘he came and did it’
- e. wɔbɛnom? ‘you will drink it’

The words in (66a-c) contain only vowels that are [-low, +ATR] and the words in (66d-e) contain only [-low, -ATR] vowels.

In traditional vowel harmony terms the presence of a [+low] vowel blocks the leftward spread of ATR, as in (67). 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. For example, the words in (67) each contain a [+low] vowel surrounded by [-low] vowels. The [+low] blocker has the same ATR feature value as the [-low] vowels to its right while those to the left have a different ATR feature value. This distribution of the harmonic feature with respect to the [+low] blocker is maintained throughout Akan but in other languages the blocker could have the same harmonic feature value as vowels to its left.

⁴Akan may include some additional words that further complicate this pattern. Dolphyne (1988) includes two counterexamples to the general Akan pattern with the stems [pinkjɛ] ‘come close’ and [njinsɛn] ‘be pregnant’. She also includes a longer word with additional morphology [onjinsɛnɪ] ‘she became pregnant’, which further illustrates the potential for multiple harmony domains without a blocker. However, I maintain the blocking analysis documented in the above citations for Akan and abstract away from these counterexamples in order to illustrate how the basic concept of blocking can be computed.

(67) Akan words with [+low] blockers (adapted from Clements 1976)

- a. pɪræko ‘pig’
- b. mɪkəkæri ‘I go and weight it’
- c. obisar ‘he asked’
- d. okog^warɪ? ‘he goes and washes’


So in (67a-b) the final vowel is [-low, +ATR], the blocker is [+low, +ATR], and the initial vowels are [-low, -ATR]. The opposite order is illustrated in (67c-d) where the final vowel and blocker are both [-ATR], but the first two vowels are [-low, +ATR]. I make no assumptions about how vowels obtain their ATR features. I only observe and account for the surface distribution of feature associations. So I do not make any claims about whether or not the [+low] blockers trigger ATR harmony to their left. The absence of words with [+ATR] [e] followed by [-ATR] [a] could be an accidental gap given the prevalence of words with [+ATR] [u], [i], or [o] followed by [-ATR] [a], but this remains speculation.

4.1.1 Blocking over Multi-tiered ARs

In this section I present my novel analysis of blocking vowels. I investigate the complexity of blocking over multi-tiered ARs using methods similar to those developed by Jardine (2016b) to analyze tone patterns. I show that over multi-tiered ARs the Akan ATR harmony pattern with blocking is local and fits into my new ASL^{VH} class of ARsets.

Over multi-tiered ARs a CNL is a set of FSCs. The CNL in (68) describes the ATR harmony generalization with blocking in Akan using only three FSCs. These constraints use only the successor relation between elements on a tier and the association relation between elements on different tiers and so they are also local.


(68) Set of FSCs for blocking in Akan



The first constraint in (68) states that no word can contain two vowels associated to the same [-low] feature which are also associated to different ATR features and this captures the basic ATR harmony generalization. Together, the second and third constraints state that no string can contain a vowel associated to a [+low] feature which is associated to a different ATR feature than the vowel that succeeds it; and this captures the fact that on the surface Akan blockers are associated to the same ATR feature as vowels to their right, regardless of whether they participate in an assimilation process or become associated to the [-ATR] feature by default.

The FSC in (68a) does not include the successor relation despite the pattern being computed over strings of vowels without intervening consonants; this omission captures the fact that no two vowels at any distance can be simultaneously associated to the same [-low] feature and different ATR features. In other words, all vowels associated to a single [-low] feature must also be associated to the same ATR feature so Akan ATR harmony is diffuse. Since it does not include a [+low] feature (68a) also allows an AR to contain a [+low] blocking vowel, as in (69). The AR in (69a) includes two different ATR features, but they are associated to [-low] vowels on either side of a [+low] vowel.


(69) Example for (68a)



Because (69a) obeys the NCC and the [+low] vowel intervenes, the vowels on either side of it are associated to different iterations of a [-low] feature. Thus the grammatical AR in (69a) does not contain the FSC in (68a) and (69b) is ungrammatical because it does.

In addition, the FSCs in (68b-c) specifically prevent the [+low] blocking vowels from being associated to an ATR feature that differs from the ATR feature associated to the [-low] vowel which succeeds it. This restriction enforces that the [+low] blocking vowels will always be associated to the same ATR feature as successive [-low] vowels rather than preceding [-low] vowels on the surface. These FSCs also allow a string to end with multiple [-low] vowels as long as the [+low] vowel is associated to the same ATR feature as them.

(70) Example for (68b-c)



In the grammatical (70a) the [+low] and the rightmost [-low] vowels are associated to a single [+ATR] feature. The ungrammatical (70b-c) both contain [+low] vowels that are associated to the same ATR feature as the leftmost [-low] vowel which violates (68b) and (68c), respectively.

The analysis of ATR harmony in Akan presented in this section predicts that longer vowel strings with more than one [+low] vowel should also be possible. The FSCs in (68) further predict that sequences of [+low] vowels do not have to be associated to the same ATR feature. This prediction is supported by only a few pieces of data from Akan found in Dolphyne (1988).


(71) Akan words with multiple [+low] vowels (Dolphyne 1988)

- a. muakæri ‘you(pl.) have weighed it’
- b. ægya ‘father’
- c. mætwa ‘I’ve cut it’
- d. ægua ‘chair’
- e. æhua ‘one who begs for food’

The words in (71a-c) contain more than one [+low] vowel with no other vowels between them. The words in (71d-e) contain a [-low, +ATR] vowel between two [+low] vowels; so the first vowel is [+low, +ATR] and the final vowel is [+low, -ATR].

This data provides crucial evidence for the necessity of the FSC in (68c). Without the addition of the ordering relation between vowels, the FSC would predict that a word with multiple [+low] vowels in the order of those found in (71a) should not be grammatical, but according to Dolphyne (1988) (71a) is a grammatical word of Akan.


(72) [muakæri] ‘you(pl.) have weighed it’




While this was the only single example of a word with vowels in this order, it suggests that such words may be possible in Akan and so my analysis should account for them. One option could have been to remove the relevant constraint from the set of FSCs entirely, but doing so would predict that ARs like (70c) are grammatical when Akan provides no evidence of words with such a sequence of vowels. The alternative was to add something to the constraint which would include ARs like (72) while still excluding (70c). I thus included the ordering relation between vowels within the constraint (68c) so that it only applies to vowels that are directly adjacent to one another. In this way the AR in (72) is still grammatical because the FSC in (68c) allows strings to include a sequence of multiple [+low] vowels in any order. I did not find additional examples with multiple [+low] vowels in the same order as in (71a)—[+ATR,+low][-ATR,+low]—but I include the ordering relation in (68b) as well to predict the same behavior. A comparison of grammatical and ungrammatical ARs from (68) and (72) are provided below to illustrate the necessity of this additional ordering relation within the FSC in (68c).

(73) Comparison of ARs using FSC from (68c)

a. ARs from (70)



b. Grammatical string from (72)



In (73a) I show that an ungrammatical AR is forbidden by the FSC in (68c). The ungrammatical AR on the left contains the substructure forbidden by the FSC in the middle, but the grammatical AR on the right does not. This example illustrates how the FSC in (68c) is still able to enforce that blocking vowels are associated to the same ATR feature as [-low] vowels to their right, but not those to their left. In addition, the FSC predicts that a word with multiple [+low] vowels in the order shown in (73b) is grammatical because it does not contain the substructure of (68c).

4.1.2 Surface diffusion is local

This dissertation focuses only on surface representations and Akan provides an example of a pattern in which vowel harmony assimilation with blocking is represented by diffuse ARs. Diffuse ARs consist of a single harmonic feature that is multiply associated to a string of successive vowels. The analysis of Akan provided here illustrates diffusion because ATR features are multiply associated to all vowels within a domain, either including and to the right of a [+low] blocker or all vowels to the left of a [+low] blocker. These diffuse ARs are local on the surface because they consist of a domain defined by a single ATR feature node, they include a contiguous span of vowels, but they are not bounded in length; or when two different ATR features are present, one succeeds the other regardless of how many vowels are associated to each.

4.1.3 Blocking over strings

In this section I present my analysis of blocking vowels which highlights the simplicity of patterns that utilize them. I investigate the complexity of blocking over strings of vowels using mathematically proven methods to test the logically possible stringset based on Akan. I show that as a stringset the Akan ATR

harmony pattern described below is provably SL_2 .

The ATR harmony generalization in Akan can be described as right-to-left harmony with [+low] blockers [a, æ] (Casali 2012; Clements 1976; Dolphyne 1988; O’Keefe 2004), but on the surface it is realized as the diffusion of a single ATR feature to all vowels within a sequence of [-low] vowels. The [+low] blocking vowels will also have the ATR feature of one [-low] sequence within a word but they require two such sequences with different ATR feature values. Previous work has argued that for a right-to-left harmony pattern, a nonfinal vowel must acquire its [+ATR] feature via harmony from an underlyingly specified [+ATR] trigger and [-ATR] features are assigned by default to any remaining vowels which are not specified underlyingly (Casali 2012; Dolphyne 1988). Further, previous analyses of blocking have claimed that blockers are neutral because they do not participate in harmony (van der Hulst 2016, 2018; van der Hulst and Smith 1986; van der Hulst and van de Weijer 1995); but the Akan pattern appears to contradict that claim on the surface. The novel contribution of my analysis is to treat [+ATR] and [-ATR] vowels equally without relying on dominance or morphological control. The FSCs posited below are thus able to describe the Akan pattern as a surface phonotactic pattern regardless of how blockers might acquire their ATR feature.

In this section I evaluate the Akan blocking pattern over strings of vowels which make up the stringset \mathbb{L}_A . The sets of strings which are and are not members of the stringset \mathbb{L}_A are briefly described in Table 8 below.

Table 8

strings $\in \mathbb{L}_A$	strings $\notin \mathbb{L}_A$
ii, i, ia, iæ, æa, ææ, ...	i, i, ai, æi, æu, au, ...
iii, iiu, iie, iio, iiæ, ...	ii, oii, eii, oii, aii, ...
iii, iio, iie, iio, ...	ii, iu, ie, io, ...
ia, iæ, æua, æaa, aaæ, aæa, æaæ, ...	

This stringset includes strings with only [+ATR] vowels, only [-ATR] vowels, or both when a [+low] vowel intervenes. It also includes strings with more than one [+low] vowel, which I use to make predictions about the types of strings that are possible in \mathbb{L}_A . These predictions are borne out in the attested words of Akan, and are discussed in more detail below.

The blocking pattern in Akan ATR harmony is SL_2 over strings because it obeys Suffix Substitution Closure (SSC) and it can be described by a Conjunction of Negative Literals (CNL) interpreted with the successor relation (\triangleleft). The CNL that describes the blocking pattern in \mathbb{L}_A consists of strings containing only two vowels as shown in (74).

(74) SL CNL for Blocking in \mathbb{L}_A

$$\begin{aligned} \mathbb{L}_A: & \neg [+ATR, -low][\neg ATR, -low] \wedge \neg [\neg ATR, -low][+ATR, -low] \wedge \neg [+ATR, +low][\neg ATR, -low] \wedge \neg \\ & [\neg ATR, +low][+ATR, -low] = \\ & \neg i\bar{i} \wedge \neg u\bar{u} \wedge \neg e\bar{e} \wedge \neg o\bar{o} \wedge \neg \bar{a}i \wedge \neg i\bar{u} \wedge \neg u\bar{u} \wedge \neg e\bar{u} \wedge \neg o\bar{u} \wedge \neg \bar{a}u \wedge \neg i\bar{e} \wedge \neg u\bar{e} \wedge \neg e\bar{e} \wedge \neg o\bar{e} \wedge \neg \bar{a}e \\ & \wedge \neg i\bar{o} \wedge \neg u\bar{o} \wedge \neg e\bar{o} \wedge \neg o\bar{o} \wedge \neg \bar{a}o \wedge \neg i\bar{i} \wedge \neg u\bar{i} \wedge \neg e\bar{i} \wedge \neg o\bar{i} \wedge \neg \bar{a}i \wedge \neg i\bar{u} \wedge \neg u\bar{u} \wedge \neg e\bar{u} \wedge \neg o\bar{u} \wedge \neg \bar{a}u \\ & \wedge \neg i\bar{e} \wedge \neg u\bar{e} \wedge \neg e\bar{e} \wedge \neg o\bar{e} \wedge \neg \bar{a}e \wedge \neg i\bar{o} \wedge \neg u\bar{o} \wedge \neg e\bar{o} \wedge \neg o\bar{o} \wedge \neg \bar{a}o \end{aligned}$$

The CNL with ATR features in (74) provides a clear and succinct statement of the ATR harmony generalization in \mathbb{L}_A , but it is only a shorthand as feature strings have not been studied on their own as a distinct representation within the FLT framework. That statement is equivalent to the CNL that lists all of the forbidden vocalic substrings of length $k=2$. In either format, the CNL for \mathbb{L}_A allows for longer strings to include the listed vowel combinations as long as another vowel intervenes. Interpreted with the successor ordering relation (74) describes \mathbb{L}_A as a SL_2 stringset and so these constraints are local. On their own a majority of the NLS in (74) could generate a grammar that differs from the grammars generated by the other NLS within the CNL, but conjoined they describe \mathbb{L}_A . However, listing an arbitrary set of segmental restrictions does not succinctly capture the intuitive generalization of ATR harmony with blocking in \mathbb{L}_A the same way that the CNL with two featural NLS does.

In this section I have shown that over strings of vowels ATR harmony with blocking in Akan fits into both the SL_2 and ASL^{VH} classes and is thus local. In addition, this Akan pattern uses diffuse multi-tiered ARs which represent locality via multiple association of the harmonic feature.

4.2 Transparent vowels: Finnish

In this section I discuss a frequently cited example of transparency with backness harmony in Finnish. I first translate actual words of Finnish into strings of vowels in order to determine that this pattern is SP_2 . I then present my novel FSC analysis of Finnish back harmony with transparency over multi-tiered ARs to show that this pattern is also ASL^{VH} .

Finnish provides an example of backness harmony with four transparent vowels. The Finnish vowel inventory in Table 9 consists of 16 vowels with contrastive length and three main featural distinctions: $[\pm \text{back}]$, $[\pm \text{low}]$, and $[\pm \text{round}]$ (Ringen and Heinamaki 1999; Välimaa-Blum 1986). The four vowels transparent to backness harmony, $[i, i:, e, e:]$, are all $[-\text{back}, -\text{round}, -\text{low}]$. Of the harmonizing vowels $[y, y:, u, u:, \emptyset, \emptyset:, o, o:]$ are all $[\text{+round}, -\text{low}]$ while $[\bar{a}, \bar{a}:, \alpha, \alpha:]$ are all $[-\text{round}, \text{+low}]$. The $[\text{+back}]$ vowels are $[u, u:, o, o:, \alpha, \alpha:]$ and the $[-\text{back}]$ vowels are $[i, i:, e, e:, y, y:, \emptyset, \emptyset:, \bar{a}, \bar{a}:]$. Finnish transparent vowels do not contrast in backness so they are all $[-\text{back}]$, but harmonizing vowels can be either $[\text{+back}]$ or $[-\text{back}]$. The

major difference between harmonizing and transparent Finnish vowels is characterized by the low and round feature values: transparent vowels are all [-low, -round] so harmonizing vowels have a positive value for either the low and/or the round feature. The vowel chart in Table 9 is adapted from Ringen and Heinamaki (1999).

Table 9: Finnish Vowels

	-round	+round	
-low	i	y	-back
	e	ø	
+low	æ		+back
	ɑ		
-low		u	
		o	

The Finnish back harmony generalization is that all of the harmonizing vowels in a root and the harmonizing suffix vowels will have the same back feature value (Nevins 2010; Ringen and Heinamaki 1999; Välimaa-Blum 1986; van der Hulst 2017). 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. Sources for the Finnish data in this section are indicated using the following abbreviations: (RH) for Ringen and Heinamaki (1999) and (VB) for Välimaa-Blum (1986).

(75) Finnish harmonizing vowels share a back feature value

- a. pøytæ ‘table’ (RH)
- b. kæntæ: ‘turn’ (VB)
- c. tykætæ ‘like’ (VB)
- d. pouta ‘fine weather’ (RH)
- e. murta: ‘break’ (VB)
- f. kokata ‘cook’ (VB)

For example, the vowels in (75a-c) are all [-back] and the vowels in (75d-f) are all [+back]. In addition, the words in (75) contain only vowels that are [+round] or [+low] but not both because the [+low, -round] vowels do not have [+round] counterparts.

Transparent vowels, however, do not block or undergo harmony so [+back] harmony appears to skip over the [-back, -round, -low] vowels [i, i:, e, e:] in Finnish. The words in (76) show that the two vowels [i] and [e] can occur anywhere within a word and all other vowels must still have the same back feature.

(76) Finnish back harmony skips over transparent vowels (VB)

- | | |
|-----------------------------|------------------------|
| a. æitiæ ‘mother’ | d. ruveta ‘start’ |
| b. ky:neltæ ‘from the nail’ | e. tuolia ‘chair’ |
| c. værikæs ‘colorful’ | f. lukea ‘read (inf.)’ |


In (76a-c) all vowels are [-back] including [i] and [e], but (76d-f) clearly illustrate the transparency of these two vowels because all other vowels are [+back] regardless of where in the word [i] or [e] occur.

4.2.1 Transparency over multi-tiered ARs

In this section, I present a new perspective on the surface representation of transparent vowels. Previous work accounts for the behavior of transparent vowels by positing that they are unspecified for backness. The novel contribution of my analysis is to treat transparent vowels in the same way as harmonizing vowels on the surface. Both harmonizing and transparent vowels obey Full Specification (FS) and are associated to a feature on each feature tier. Multi-tiered ARs also allow us to analyze both blocking in Akan and transparency in Finnish with constraints that use the same relations. Over multi-tiered ARs both patterns are describable with FSCs that use the successor relation and so both patterns can be classified as ASL^{VH}. In this subsection I propose a Forbidden Substructure Grammar (FSG) over multi-tiered ARs. I do this in order to argue that Finnish back harmony can be captured by a set of phonotactic restrictions over multi-tiered ARs without the need for underspecification of transparent vowels on the surface.


If you recall the data in (75) and (76) showed that in a Finnish word all vowels must have the same back feature except [i] and [e], which are [-back, -low, -round]. 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 preceding back feature. More specifically, when a vowel associated to either [+round] or [+low] is also associated to [-back] the four FSCs in (77) forbid that [-back] feature from either succeeding or being succeeded by a [+back] feature. In short, these four FSCs describe the fact that Finnish harmonizing vowels have a positive value for either the round or the low feature.

(77) Finnish FSCs



The ordering relation on the back tier is crucial to allow the transparency of certain [-back] vowels. The ARs in (78) illustrate how the Finnish FSCs rule out ungrammatical disharmonic words. The AR for the grammatical Finnish word [pouta] ‘fine weather’, shown in (78a), contains both a [+round] and a [+low] non-initial vowel as well as a single [+back] feature, which demonstrates fully diffuse back harmony. ARs with fully diffuse back harmony are also grammatical in a language like Finnish for words that do not contain any transparent vowels.

(78) [pouta] ‘fine weather’ (RH)




The hypothetical Finnish word, [poutæ] in (78b), however, contains the forbidden substructure of (77a) in bold and red. In (78b) the final vowel does not harmonize with the penultimate vowel because they are associated to different back features.

Enforcing full specification allows surface ARs of Finnish words to represent back harmony iteratively across transparent vowels. For example, the words in (76) 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 (79). 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 preceded or succeeded by a [+back] feature. This restriction enforces [+back] assimilation across transparent vowels, i.e. iterative assimilation. 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, -round]. So, in words with more than one back feature any [-back] vowel must be transparent and all other vowels must be [+back].

(79) [ruveta] ‘start’ (VB)

a.



b.




In (79a) the [u] and [a] 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 (79a) 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 (77). The AR in (79b), on the other hand, contains a [-back, +low] vowel, and so (77a) 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 assimilate in backness and the same FSCs that capture Finnish back harmony in (79a) also enforce assimilation across a transparent vowel by marking words like (79b) as ungrammatical.

In summary, the iterative back harmony pattern with transparent vowels in Finnish can be described using four local FSCs over multi-tiered ARs and so it fits into the ASL^{VH} class of ARsets. The Finnish pattern is thus considered local because the FSCs that describe the pattern are connected by the successor ordering relation between elements on a tier and the association relation between elements on different tiers.

4.2.2 Surface iteration can be local

The analysis of Finnish provided here demonstrates that iterated ARs can be local. Here locality means that iterated 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. In Finnish there are grammatical words like (78a) which are represented by an AR with diffusion on the back tier, but there are also words like (79a) which contain two [+back] harmonizing vowels with one or more transparent vowels between them. 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. Such ARs are made possible by the lack of backness contrast among Finnish transparent vowels and the autosegmental assumptions I adopted and discussed in Chapter 3. The NCC prevents Finnish from using only diffuse ARs with transparent vowels 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, and so Finnish ARs do not violate the OCP. The assimilation of [+back] between two vowels in words like (79a) is different from diffusion because there are two [+back] features and they 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 an intervening [-back] feature. I call the surface representation of this other type of assimilation *iterated*. Iteration is represented as a surface AR in which two non-successive features on a tier have the same value and the intervening feature on that tier has the opposite value, as shown in (79) and (80). The Finnish word in (80) is taken from Sulkala and Karjalainen (1992) page 51, example 218 which I abbreviate as (SK).

(80) [maisemia] ‘scenery.plural.partitive’ (SK)



In (80) the [-back] feature is diffused so it is associated to multiple transparent vowels, but the [+back] feature is represented in multiple iterations. This surface configuration allows the AR to represent [+back] harmony locally because the two iterations of [+back] are adjacent to only a single [-back] feature regardless of how many transparent vowels it associates to.

4.2.3 Transparency over strings

Over strings, it is clear that blocking in Akan and transparency in Finnish must be treated differently. Akan is SL_2 and so its FSCs over strings are interpreted with the successor ordering relation. In this subsection I investigate the complexity of vowel harmony with transparency over strings and show that it is SP_2 so its FSCs over strings are interpreted with the precedence ordering relation.

The Finnish back harmony generalization is described as left-to-right harmony with [-back, -round, -low] transparent vowels [i, e] (Nevins 2010; Ringen and Heinamaki 1999; Välimaa-Blum 1986; van der Hulst 2017), but on the surface it is realized as the sharing of a back feature between harmonizing vowels in roots and suffixes that could be separated by a sequence of transparent vowels of any length. I evaluate the Finnish transparency pattern over strings of vowels which make up the stringset \mathbb{L}_F in order to precisely determine that \mathbb{L}_F is not SL but fits into the class of SP_2 stringsets. The sets of strings which are and are not members of \mathbb{L}_F are briefly summarized in Table 10 below.

Table 10

strings $\in \mathbb{L}_F$	strings $\notin \mathbb{L}_F$
yø, oɑ, øæ, ...	yu, øo, æɑ, ...
iu, eo, ai, ...	æu, oy, aø, ...
øæy, yiæ, æeø, ...	yoø, æuo, ayæ, ...
uoɑ, oia, æeu, ...	yea, æiu, aeø, ...
yiiiæ, æeieio, æeeeu, ...	yiiiɑ, æieieo, æeeeo

The transparency pattern in Finnish back harmony is not SL_k for any k because it does not obey SSC, as shown in Table 11.

Table 11: Suffix Substitution Closure test for \mathbb{L}_F

$x = e^{k-1}$	$wxy = ye^{k-1}æ$	$\in \mathbb{L}_F$
	$vxz = ue^{k-1}a$	$\in \mathbb{L}_F$
	$wxz = ye^{k-1}a$	$\notin \mathbb{L}_F$

In other words, the restriction on the backness of harmonizing vowels in Finnish holds regardless of how many transparent vowels might intervene. Exchanging the suffixes from two strings with harmony in different back features within \mathbb{L}_F yields two strings that are not in \mathbb{L}_F because the harmonic feature no longer iterates across the transparent vowel(s). The test in Table 11 thus illustrates that back harmony holds across any arbitrary distance so it cannot be SL_k for any k .

The Finnish back harmony pattern with transparency is Strictly 2-Piecewise (SP_2) because it obeys Subsequence Closure (SC) and is describable with a CNL over strings of vowels with length $k = 2$ that is interpreted with the precedence ordering relation ($<$). The CNL forbids words from including two vowels at

any distance that are [+low] or [+round] and have different back feature values, as shown in (81). The NLs in (81) make up exactly the set of strings—whether featural strings or vocalic strings—of length $k = 2$ that are not members of \mathbb{L}_F . Since the NLs have length $k = 2$, \mathbb{L}_F is SP_2 . Because these NLs are interpreted with precedence they hold regardless of how many transparent vowels might occur between the two harmonizing vowels stated in the restriction and thus would account for longer strings like those in the last two rows of Table 10 as well.

(81) CNL for \mathbb{L}_F transparency pattern, interpreted with precedence ($<$)

$$\begin{aligned} \mathbb{L}_F: & \neg [-\text{back}, +\text{round}][+\text{back}, +\text{round}] \wedge \neg [+ \text{back}, +\text{round}][-\text{back}, +\text{round}] \\ & \wedge \neg [-\text{back}, +\text{low}][+\text{back}, +\text{low}] \wedge \neg [+ \text{back}, +\text{low}][-\text{back}, +\text{low}] \wedge \neg [-\text{back}, +\text{round}][+\text{back}, +\text{low}] \wedge \neg \\ & [+ \text{back}, +\text{round}][-\text{back}, +\text{low}] \wedge \neg [-\text{back}, +\text{low}][+\text{back}, +\text{round}] \\ & \wedge \neg [+ \text{back}, +\text{low}][-\text{back}, +\text{round}] = \\ & \neg \text{yu} \wedge \neg \text{uy} \wedge \neg \text{yo} \wedge \neg \text{oy} \wedge \neg \text{ya} \wedge \neg \text{ay} \wedge \neg \text{øu} \wedge \neg \text{uø} \wedge \neg \text{øo} \wedge \neg \text{oø} \wedge \neg \text{øa} \wedge \neg \text{aø} \wedge \neg \text{æu} \wedge \neg \text{uæ} \wedge \neg \\ & \text{æo} \wedge \neg \text{oæ} \wedge \neg \text{æa} \wedge \neg \text{aæ} \end{aligned}$$


As with the CNL for \mathbb{L}_A in (74), a majority of the NLs in (81) could generate a separate grammar from those generated by the other NLs within the CNL, but conjoined they accurately describe \mathbb{L}_F . Again, however, listing an arbitrary set of restrictions on unanalyzed symbols does not capture the intuitive generalization of back harmony with transparency in \mathbb{L}_F . Unlike \mathbb{L}_A , the featural CNL for \mathbb{L}_F is not much more succinct than the vocalic CNL but the two are still equivalent. Remember that the featural strings here are solely a shorthand to describe the pattern, although for Finnish the featural CNL is not much shorter than the vocalic one.

A multisyllabic word provides a clear example of how the CNLs for \mathbb{L}_F work over any distance because they are interpreted with precedence. For example, The vowel string on the left in (82a) is taken directly from the grammatical Finnish word [maisemia] ‘scenery.plural.partitive’ which has only an initial and a final [+low] vowel. The NLs on the right restrict the back features values of only two [+low] vowels. Interpreting the NLs with precedence means that the restrictions on these [+low] vowels hold regardless of how many transparent vowels intervene between them. So the vowel string on the left is grammatical because it does not contain any of the forbidden pieces of a string (substrings) on the right.

(82) a. [maisemia] ‘scenery.plural.partitive’ (SK)

$$\begin{array}{ccc} \text{a} \longrightarrow \text{i} \longrightarrow \text{e} \longrightarrow \text{i} \longrightarrow \text{a} & & \neg [-\text{back}, +\text{low}][+\text{back}, +\text{low}] \wedge \\ & \curvearrowright & \neg [+ \text{back}, +\text{low}][-\text{back}, +\text{low}] = \\ & < & \neg \text{æa} \wedge \neg \text{aæ} \end{array}$$

b. Ungrammatical vowel string



$$\neg [-\text{back}, +\text{low}][+\text{back}, +\text{low}] = \neg \text{æɑ}$$

If the first or last vowel were changed to have a different back feature value then the string would violate one of the constraints on [+low] vowels and be ungrammatical. In (82b) the initial vowel is [-back,+low] [æ] and the final vowel is [+back,+low] [ɑ]. This combination of vowels is forbidden by the NLs on the right regardless of how many transparent vowels intervene so this vowel string is ungrammatical. The visual representation of the precedence relation between the [+low] vowels in (82) with a curved arrow makes it clear why vowel harmony with transparency is considered a long distance pattern even over strings of vowels; again, the constraints on the two [+low] vowels hold regardless how many transparent vowels intervene. Over strings, allowing any arbitrary number of vowels to intervene makes SP_k constraints non-local.

Finnish back harmony with transparent vowels is also Tier-based strictly 2-local (TSL_2). The CNL in (81) can also be used for a TSL analysis of Finnish when it is interpreted with the successor relation over a string tier which excludes the transparent vowels (Heinz, Rawal, and Tanner 2011).

(83) TSL grammar for \mathbb{L}_F transparency pattern

a. Tier of vowels excluding transparent [-back, -high, -round]: $T = y, \emptyset, \text{æ}, \text{ɑ}, \text{u}, \text{o}$

b. CNL interpreted with successor (\triangleleft) over T :

$$\begin{aligned} & \neg [-\text{back}, +\text{round}][+\text{back}, +\text{round}] \wedge \neg [+ \text{back}, +\text{round}][-\text{back}, +\text{round}] \\ & \wedge \neg [-\text{back}, +\text{low}][+\text{back}, +\text{low}] \wedge \neg [+ \text{back}, +\text{low}][-\text{back}, +\text{low}] \wedge \neg [-\text{back}, +\text{round}][+\text{back}, +\text{low}] \\ & \wedge \neg [+ \text{back}, +\text{round}][-\text{back}, +\text{low}] \wedge \neg [-\text{back}, +\text{low}][+\text{back}, +\text{round}] \\ & \wedge \neg [+ \text{back}, +\text{low}][-\text{back}, +\text{round}] = \\ & \neg \text{yu} \wedge \neg \text{uy} \wedge \neg \text{yo} \wedge \neg \text{oy} \wedge \neg \text{ya} \wedge \neg \text{ay} \wedge \neg \text{øu} \wedge \neg \text{uø} \wedge \neg \text{øo} \wedge \neg \text{oø} \wedge \neg \text{øɑ} \wedge \neg \text{ɑø} \wedge \neg \text{æu} \wedge \neg \\ & \text{uæ} \wedge \neg \text{æo} \wedge \neg \text{oæ} \wedge \neg \text{æɑ} \wedge \neg \text{ɑæ} \end{aligned}$$

Excluding the transparent vowels from the tier projection allows Finnish to be computed locally because the harmonizing vowels are adjacent on the tier regardless of how many transparent vowels might intervene. Even with a string of vowels the TSL grammar in (83) allows the same long-distance pattern—which otherwise required SP—to be described in a local way.

(84) a. [maisemia] ‘scenery.plural.partitive’ (SK)



b. Ungrammatical vowel string



Again the combination of a [+back] vowel followed by a [-back] vowel as in (84b) is ungrammatical and it violates the NL on the right, as shown in bold and red. However, this time the NL is interpreted locally over the tier projection with the successor relation.

In this section I have shown that over strings of vowels back harmony with transparency in Finnish falls into the SP_2 , TSL_2 , and ASL^{VH} classes and is thus local over a tier projection which excludes the transparent vowels and over multi-tiered ARs. This autosegmental locality is a direct result of the fact that Finnish transparent vowels do not contrast in backness. In addition, this pattern uses iterative multi-tiered ARs which represent locality via multiple association of a single feature which intervenes between two iterations of the harmonic feature.

4.3 Conclusion

In this chapter I have shown that vowel harmony patterns with neutral vowels can be described using surface phonotactic constraints called FSCs over both strings of vowels and multi-tiered ARs. ATR harmony with blocking in Akan even with an ATR contrast between blocking vowels is SL_2 and ASL^{VH} . This means that Akan is local over both strings and multi-tiered ARs. On the other hand, back harmony with transparency and no neutral vowel contrast in Finnish is SP_2 and ASL^{VH} even without underspecification. This means Finnish is only local over multi-tiered ARs and not strings. While these two vowel harmony patterns fit into different subregular stringset classes both are ASL^{VH} . So the ASL^{VH} class crosscuts the subregular stringset hierarchy and includes patterns that fit into more than one stringset class.

This chapter also shows that there are clear computational differences between blocking and transparency in vowel harmony. The SL and SP stringset classes are not comparable. Even over multi-tiered ARs the surface representations of Akan and Finnish utilize different harmony mechanisms: diffusion and iteration,

respectively. It is difficult to conceive of a pattern with non-local diffusion but Finnish transparency is local over multi-tiered ARs specifically because Finnish does not have a harmonic contrast between its transparent vowels. Neutral vowel contrasts thus play a significant role in the complexity of iterative transparency but not diffuse blocking. The computational significance of these differences between blocking and transparency will be explored even further with more complex patterns in the next chapter.

5 ARs reduce the complexity of neutral vowels

In this chapter I analyze two vowel harmony patterns with neutral vowels which have previously been analyzed as more complex than SL or SP over strings. I show that different types of neutral vowels have different effects on the computational complexity of vowel harmony patterns. For example, transparent vowels generally make a vowel harmony pattern more complex than blocking vowels do. In the previous chapter I showed that the Akan ATR harmony pattern with blocking vowels is SL_2 over strings and the Finnish back harmony pattern with transparent vowels is SP_2 over strings. Similarly, in this chapter I show that the ATR harmony pattern with blocking vowels in Tutrugbu is Locally Testable (LT) over strings and ASL^{VH} , but the Eastern Meadow Mari (EMM) back harmony pattern with transparent vowels is LT over strings and is more complex over ARs than patterns which fit into the ASL^{VH} class.

So far work on the complexity of Tutrugbu has focused on comparing it as a function to other phonological functions and it seems to stand out; but as a phonotactic pattern Tutrugbu fits into the same class as many other vowel harmony patterns. Recent work on Tutrugbu in McCollum et al. (2020) analyzes its back harmony pattern as an input-output mapping described by a function which is more complex than the (sub)sequential class of functions which has otherwise been argued to include most phonological processes (Chandlee 2014; Heinz and Lai 2013). My analysis of Tutrugbu as a surface phonotactic pattern reveals that over multi-tiered ARs Tutrugbu is more like other phonotactic vowel harmony patterns than other phonological processes since it also fits into the ASL^{VH} class of ARsets.

Eastern Meadow Mari (EMM) has previously been analyzed as an example of a vowel harmony pattern with a positional restriction but there is more to it. I show that EMM differs from other vowel harmony patterns because it utilizes positional transparency and the transparent vowels contrast in the harmonic feature. The harmonic contrast between transparent vowels causes the EMM back harmony pattern to be more complex than other vowel harmony patterns over both vowel strings and multi-tiered ARs. This complexity suggests the definition of transparency must exclude harmonic contrasts in order to maintain the hypothesis that all vowel harmony patterns are ASL^{VH} .

In section 5.1 I expand the ASL^{VH} class to include Tutrugbu and in 5.2 I show that Eastern Meadow Mari (EMM) does not fit into the ASL^{VH} class.

5.1 Tutrugbu

Tutrugbu is a language of the Kwa family which is spoken in the Volta region of Ghana. Tutrugbu utilizes a phonetic 7-vowel system that is historically reduced from a 9-vowel system (Essegbey 2019; McCollum and Essegbey 2020; McCollum et al. 2020). The crucial distinctions for ATR harmony are between the [+ATR] [i,

u, e, o] and [-ATR] [ɛ, ɔ, a] vowels as well as between the [+high] [i, u] and [-high] [e, o, ɛ, ɔ, a] vowels. The [+low] [a] vowel also plays an important role for ATR harmony in some contexts, but more on that later.

The important point to note is that based on acoustic evidence, the [+high] vowels do not contrast in ATR because the Tutrugbu 7 vowel system is crucially missing the two [+high, -ATR] vowels [ɪ, ʊ]. However, Tutrugbu does contain a historical trace of the [+high, -ATR] vowels in the phonological behavior of [-high, -ATR] mid vowels [ɛ, ɔ] (Essegbey 2019; McCollum and Essegbey 2020). Morphemes that assimilate to [+ATR] with [i] will assimilate to [-ATR] with ɛ and morphemes that assimilate to [+ATR] with [u] will assimilate to [-ATR] with [ɔ]. McCollum and Essegbey (2020) maintain the transcriptions of these [-ATR] vowels as [-high] mid vowels [ɛ, ɔ] and add a superscript to denote their phonological behavior as contrasting with [+high] vowels, but McCollum et al. (2020) use the [+high, -ATR] symbols [ɪ, ʊ] to transcribe them.

Table 12: Tutrugbu Vowels

	+ATR	-ATR	
+high	i	(ɪ)	-low
	u	(ʊ)	
-high	e	ɛ	
	o	ɔ	
		a	

In order to clarify the phonological pattern on the surface I follow McCollum et al. (2020) in transcribing 9 different Tutrugbu vowels, which reflects the phonological ATR contrast in [+high] vowels even if there is no acoustic difference between [ɪ, ɛ] or [ʊ, ɔ].

The use of different transcription systems across Tutrugbu resources can cause some confusion about the data. For example, Essegbey (2019) and McCollum and Essegbey (2020) use the seven vowel transcription system that excludes [ɪ, ʊ], so their data include the words in the left column of (85) which appear to illustrate iterative harmony, as in (86a). However, in this dissertation I adopt the nine vowel transcription system used in McCollum et al. (2020) and include the [+high, -ATR] vowels [ɪ, ʊ]. The large plus symbol (+) represents a morpheme boundary; roots are the rightmost morpheme and each morpheme to the left of a root is a prefix. ATR harmony applies across morphemes and so the morpheme boundary symbol is not necessary within ARs. In this section data sources are notated using the following abbreviations: (E) for Essegbey (2019), (ME) for McCollum and Essegbey (2020), and (Mea) for McCollum et al. (2020).

(85) Tutrugbu verbs with different transcriptions

Seven vowel system


- a. bu+ka+ti (E)
- b. bu+ka+di+tseɖe (E)
- c. bu+ba+wu (ME)
- d. bu+ba+fe (ME)

Nine vowel system (Mea)


- a. bu+ka+ti ‘We do not know yet’
- b. bu+ka+di+tseɖe ‘We have not gone to tell X yet’
- c. bu+ba+wu ‘1pl+Fut+climb’
- d. bu+ba+fe ‘1sg+Fut+grow’

(86) AR of [bukati] vs. [bʊkati]

- a. [bu+ka+ti]



- b. [bʊ+ka+ti]



The examples in the right column of (85) reveal a normal circumambient blocking pattern, represented in (86b). This means that the Tutrugbu data from these three sources does not include any words with a [+ATR] feature succeeded by a [-ATR] feature. The following sections discuss the Tutrugbu ATR harmony pattern with blocking in more detail and explain why this pattern amongst ATR features is important.

5.1.1 ATR Harmony

The basic ATR harmony generalization in Tutrugbu can be described as the diffusion of ATR across morpheme boundaries with unbounded circumambient (UCA) blocking.

(87) Tutrugbu nouns with full ATR harmony (Mea)

<u>[+ATR]</u>	<u>[-ATR]</u>
a. e+bu ‘CL1+dog’	e. a+ɲɪ ‘CL1+man’
b. o+pete ‘CL3+vulture’	f. ɔ+da ‘CL3+copper’
c. i+pete ‘CL4+vulture’	g. ɪ+da ‘CL4+copper’
d. bu+ju ‘CL8+war’	h. bu+wɪ ‘CL8+axe’

(88) Tutrugbu verbs with full ATR harmony (E)

<u>[+ATR]</u>	<u>[-ATR]</u>
a. i+gi+tseɖe ‘I did not tell’	d. ɛ+ge+za ‘I did not stay’
b. o+bo+nyi ‘You will know’	e. ɔ+bɔ+ba ‘You will come’
c. be+pu ‘3pl+puncture’	f. a+ba+dɛmɔ+ɛ ‘He will go and see him’

All the words in (87)-(88) illustrate full harmony both within a polysyllabic root and across multiple morphemes but Essegbey (2019) states that “Polysyllabic roots are... often disharmonic”⁵ so I distinguish between morphemes in order to accurately represent the domain of ATR harmony in Tutrugbu. The large plus symbol (+) represents a morpheme boundary; roots are the rightmost morpheme and each morpheme to the left of a root is a prefix. Because ATR harmony applies across morphemes the morpheme boundary symbol is not needed within ARs.

In addition, the ATR harmony pattern in Tutrugbu is complicated by a very specific environment in which blocking occurs. Without a [+high] vowel on the left, a [+ATR] vowel on the right, and a [+low] vowel between them a Tutrugbu word must exhibit full harmony, as in (87)-(88) above and (89) below.

(89) Full [+ATR] harmony, no blocking (Mea)

e+ti+ke+e+be+be+wu ‘3sg+NEG+Pfv+Prog+Vent+Vent+climb’

When there is a [+high] vowel on the left and a [+ATR] vowel on the right then all the vowels to the right of a [+low] vowel are [+ATR] and all the vowels to the left of the [+low] vowel are [-ATR]; in short, the [+low] vowel requires that different ATR feature values be associated to the vowels on either side of it as in (90) below.


⁵Among the data presented in section 2.3.1.1 of Essegbey (2019), he includes only three disharmonic nouns: [kitakpu] ‘head’, [kahɔlitsa] ‘chameleon’, and [kesugba] ‘earthenware, plate’.

- (90) UCA blocking of [+ATR] harmony (Mea)
- $i+ba+wu$ ‘1sg+Fut+climb’
 - $i+t+ka+wu$ ‘1sg+Neg+Pfv+climb’
 - $i+t+ka+a+wu$ ‘1sg+Neg+Pfv+Prog+climb’
 - $i+t+ka+a+ba+wu$ ‘1sg+Neg+Pfv+Prog+Vent+climb’
 - $i+t+ka+a+ba+ba+wu$ ‘1sg+Neg+Pfv+Prog+Vent+Vent+climb’

McCollum et al. (2020) stated that “when the initial prefix vowel is [+high]... harmony is blocked by [-high] vowels.” In fact, the [+low] [a] vowel is the only vowel that blocks [+ATR] harmony when the initial vowel is [+high]. Patterns like this are called unbounded circumambient (UCA) because the environment for blocking depends upon the existence of vowels on each side of the blocker—i.e. [+high] to the left and [+ATR] to the right—regardless of the distance between them. In other words a UCA blocking pattern is one in which the blocker must be surrounded by vowels associated to specific features and those vowels can be any distance from the blocker.

Unlike the unbounded circumambience described for tone patterns in Jardine (2016a) and Jardine (2016b), Tutrugbu blocking relies on the use of multiple different features within a word; [+high], [+ATR], and [+low]. While the number of features represented in the set of Tutrugbu FSCs may increase this does not appear to affect the complexity of the pattern as compared to tone. A clear example of the unboundedness of blocking in Tutrugbu is derived from the verb ‘climb’, shown in (90) above. The examples in (90) clearly illustrate surface blocking because each word contains both a [-ATR] and a [+ATR] vowel on either side of a [+low] blocker. The second vowel in (90b-e) is also [-ATR] which demonstrates how the distinct ATR features are diffused on either side of the [+low] blocker. The same feature distribution is found regardless of how many [+low] vowels intervene which illustrates the unboundedness of blocking in Tutrugbu. All of these properties of UCA blocking in Tutrugbu are illustrated in (91).

- (91) Multi-tiered AR of UCA blocking in Tutrugbu, (90e)



While the [-ATR] feature is diffused across a majority of the vowels in (91) the UCA environment means that

it is also succeeded by a [+ATR] feature. Blocking is unbounded because the AR contains more than one feature on the ATR tier regardless of the length of the vowel string. The pattern is local over ARs because there are only two features on the ATR tier regardless of the length of the vowel string. In addition, the example in (89) above shows full harmony when the UCA environment is absent; because the first vowel is [-high] only a single [+ATR] feature is diffused across the entire word. With multi-tiered ARs the Tutrugbu blocking pattern requires the [+low] blocker to be surrounded by features on different tiers. This requirement highlights the interaction of features via their associations to vowels.

Recall that the Tutrugbu blocking pattern is more complex than many other phonological processes and so it does not fit into the (sub)sequential class of input-output mapping functions which can describe much of phonology (Chandlee 2014; Heinz and Lai 2013; McCollum et al. 2020), but when analyzed as a phonotactic pattern Tutrugbu is more like other vowel harmony patterns. In the next subsection I present my analysis of Tutrugbu ATR harmony with UCA blocking over multi-tiered ARs which shows how an enhanced representation allows us to describe the Tutrugbu generalization more simply such that it falls into the same ASL^{VH} class as many other phonological patterns. In the following subsection I present my analysis of Tutrugbu ATR harmony with UCA blocking over strings of vowels which shows that it is more complex than patterns which fit into the SL_k or SP_k stringset classes.

5.1.2 Unbounded circumambience (UCA) over ARs


The level of logic required to describe a pattern can be used to determine its complexity. Over strings of vowels the description of the UCA blocking pattern in Tutrugbu requires propositional logic and a full string analysis will be presented later. Previous work argues that propositional logic overgenerates for tone patterns so a more restrictive theory would enrich the representation instead (Heinz 2018; Jardine 2019; Lai 2015). The existence of a grammar which can generate a pattern also serves as proof of the complexity of that pattern. In this subsection I show that the Tutrugbu blocking pattern can be described with a CNL of three FSCs over the enriched representation of multi-tiered ARs which proves that the Tutrugbu pattern fits into the ASL^{VH} class of ARsets.

In order to explicitly represent the assumption that each element on the vocalic tier represents a vowel in a separate morpheme I use the morpheme boundaries which clarified the data in (87)-(89). The morpheme boundaries are represented only on the vocalic tier because they do not interact with harmony between features. Blocking in Tutrugbu is only possible if the blocker is surrounded by both a [+ATR] feature and an *initial* [+high] vowel so the ARs also include an initial word boundary represented as ‘#’. Unlike the morpheme boundaries, these word boundary symbols are crucial to the circumambient blocking pattern in Tutrugbu which


requires the initial vowel to be [+high]; so, the word boundary symbols are also represented on the high feature tier, as in (92) below.

(92) Tutrugbu words with full ATR harmony

a. [i+pete] ‘CL4+vulture’ (Mea)




b. [ɔ+bo+ba] ‘You will come’ (E)



In (92) full harmony is represented by multiple association of a single ATR feature. In (92a) the initial vowel is associated to a [+high] feature preceded by a word boundary and the final vowel is associated to a [+ATR] feature which creates the circumambient environment for blocking but there is no [+low] vowel to block harmony. In (92b) the initial vowel is associated to a [-high] feature so the environment for blocking is not present and a single [-ATR] feature is associated to each vowel.

The Tutrugbu blocking pattern is describable with a CNL of four FSCs over multi-tiered ARs, shown in (93) below. The FSCs in (93a) resembles the first Akan FSC because it forbids disharmony between two [-low] vowels. This FSC ensures that the ATR tier contains only a single feature unless a vowel is also associated to [+low]; in other words, only [+low] vowels block ATR harmony. The FSC in (93a) accomplishes this by enforcing that a [+ATR] feature cannot be preceded by a [-ATR] feature when both are associated to vowels which are also associated to a single [-low] feature. The absence of a successor relation between the ATR features in this FSC means that a [-ATR] feature also cannot be succeeded by a [+ATR] feature when both are associated to [-low] vowels.

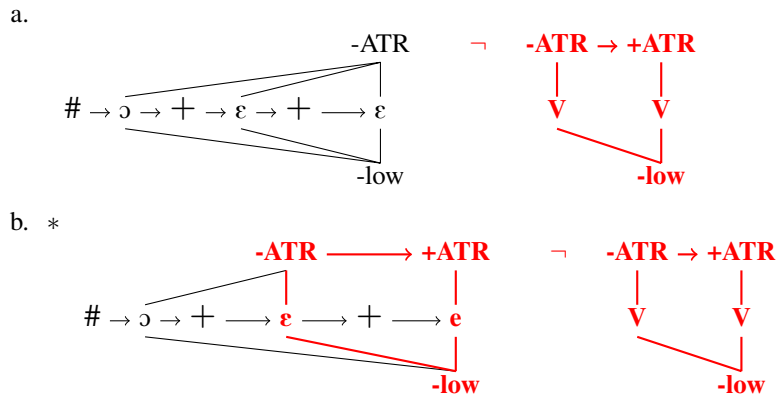
(93) ASL FSCs of Tutrugbu harmony



One example of the work (93a) does can be seen in the inflected verb [bɔzɛɛ] ‘1sg+rep+attend’. The AR of the grammatical verb in (94a) illustrates full harmony via multiple association of [-ATR] and so it does not


contain the FSC from (93a). However, if the middle vowel were able to block [+ATR] harmony then the final vowel would be associated to [+ATR]. The AR in (94b) thus contains the FSC from (93a) because the blocker and the final [+ATR] vowel are both associated to a single [-low] vowel but different ATR features, shown in bold and red.

(94) [bɔ+zɛ+dɛ] ‘1sg+rep+attend’ (E)



The FSC in (93a) also ensures that a [-ATR] feature is not preceded by a [+ATR] feature if both are associated to vowels which are in turn associated to a single [-low] vowel. Again the grammatical verb illustrates full [-ATR] harmony in (95a) and does not contain the FSC from (93a).

(95) [bɔ+zɛ+dɛ] ‘1sg+rep+attend’ (E)




If the middle vowel were somehow able to block [-ATR] harmony then the initial vowel would be associated to [+ATR]. The AR in (95b) would thus contain the FSC in (93a) because the initial vowel and the middle vowel are both associated to a single [-low] feature but different feature values on the ATR tier.


The FSCs in (93c-d) forbid disharmony on the ATR tier when the vowel succeeding a word boundary symbol is also associated to [-high]. In other words, when the first vowel is associated to [-high] the AR must

have full harmony.⁶ Over multi-tiered ARs the FSC in (93c) describes a pattern which requires propositional logic with a conditional relation over strings. The AR in (96a) of the inflected verb [ɪbaʃe] ‘1sg+Fut+grow’ illustrates the grammatical surface result of blocking [+ATR] harmony in Tutrugbu. Crucially, it does not contain the FSC of (93c) because the vowel succeeding a word boundary is also associated to [+high]. Along with the [+ATR] vowel at the right edge the initial [+high] vowel creates the circumambient environment for blocking and so the presence of two different ATR features is grammatical.

(96) a. [ɪbaʃe] ‘1sg+Fut+grow’ (Mea)




b. *




However, if the vowel succeeding the word boundary is associated to [-high] as in (96b) then the circumambient environment no longer exists and blocking is ungrammatical. The ungrammaticality of (96b) is represented by a violation of the FSC in (93c), shown in bold and red. Similarly, if the vowel succeeding the word boundary is associated to [-high] and [+ATR] then the circumambient environment no longer exists and instead of blocking (97b) illustrates iterative harmony.

(97) [ɪbaʃe] ‘1sg+Fut+grow’ (Mea)

a.



b. *




⁶At first glance the FSCs in (93) appear to also forbid words with disharmonic roots like [ka+hɔlitsa] ‘chameleon’, and [ke+sugba] ‘earthenware, plate’, but disharmonic roots are an exception to the Tutrugbu generalization described thusfar. The FSCs presented here describe a harmony pattern which applies to a domain including only root-initial vowels and prefixes so full roots are invisible to this pattern. I abstract away from strict cyclicity effects here (Chomsky and Halle 1968, Kean 1974, Levergood 1984)

Once again the AR in (97a) does not violate the FSC in (93d) which ensures that Tutrugbu does not include a [+ATR] feature succeeded by a [-ATR] feature with [-high] vowels at the beginning of a word. The AR in (97b) violates the FSC in (93d) because it contains a [+ATR] feature associated to a [-high] vowel that succeeds a word boundary and the [+ATR] feature is succeeded by a [-ATR] feature.


Longer Tutrugbu words illustrate the unboundedness of the circumambient blocking pattern enforced by the FSCs in (93). For example, the word [etikeebewu] ‘3sg+NEG+Pfv+Prog+Vent+Vent+climb’ from (89) is grammatical in Tutrugbu because it begins with a [-high] vowel and the AR contains a single multiply associated [+ATR] feature. Since there is only one ATR feature that is diffused across every vowel, (98a) does not contain the FSC in (93c).

(98) [e+ti+ke+e+be+be+wu] ‘3sg+NEG+Pfv+Prog+Vent+Vent+climb’ (Mea)




However, (98b) shows the same AR with [+ATR] blocked so there are two different features on the ATR tier while the initial vowel is associated to [-high]. This hypothetical AR is ungrammatical because it illustrates blocking, but with no initial [+high] vowel; thus it contains the FSC of (93c) as shown in bold and red. Unlike in (98b), the AR of [itikaawu] ‘1sg+Neg+Pfv+Prog+climb’ in (99a) does not contain the FSC of (93c) because the vowel that succeeds the word boundary symbol is associated to [+high].

(99) a. [ɪ+ɪ+ka+a+wu] ‘1sg+Neg+Pfv+Prog+climb’ (Mea)



b. *



The AR in (99a) clearly illustrates the circumambience of features on different tiers that is needed for blocking to occur in Tutrugbu. The [+ATR] feature is associated to a vowel that succeeds a [+low] blocking [a] and the [+high] feature succeeds a word boundary while also being associated to a vowel succeeded by a blocking [a]. In short, while both (99a) and (99b) include multiple ATR features, the initial [+high] vowel and the [+low] blocker make that grammatical in Tutrugbu.

5.1.3 UCA over strings

The Tutrugbu ATR harmony pattern can be analyzed as a set of strings made up of the nine vowels in the Tutrugbu inventory. The phonotactic approach to vowel harmony only determines which strings are and which are not included in the surface forms of a Tutrugbu stringset \mathbb{L}_T , summarized in Table 13 below. Since Tutrugbu includes disharmonic roots the strings in this analysis must only consist of vowels within the harmony domain; so to unclutter these strings I remove the morpheme boundary symbols, but the analysis crucially relies on the assumption that each vowel in one of these strings is analogous to the vowel of a separate morpheme. In short, the rightmost vowel in these strings represents the first vowel of a root and each preceding vowel within the string represents a prefix vowel.

Table 13

	strings $\in \mathbb{L}_T$	strings $\notin \mathbb{L}_T$
	iu, eo, iu, ae	iü, eö, iu, ae
Harmony:	ouo, iuö	ouü, iua
	iiee, eeee	iiee, ööö
Blocking:	iae, uao	iai, uau
	uaai, iau	eaau, ööai

The Tutrugbu blocking pattern cannot be SL_k for any k because it violates Suffix Substitution Closure (SSC). The test in Table 14 illustrates the unbounded nature of the circumambient blocking pattern in \mathbb{L}_T . Exchanging the suffixes between one string in \mathbb{L}_T that has blocking and another string that has full [-ATR] harmony and an initial [-high] vowel results in one string that is in \mathbb{L}_T and one that is not. The resulting string that is in \mathbb{L}_T has full [-ATR] harmony and the resulting string that is not in \mathbb{L}_T has [+ATR] blocked, but regardless of how many [+low] blocking vowels occur the string does not contain an initial [+high] vowel so it would not be a grammatical string of vowels in Tutrugbu.

Table 14: Suffix Substitution Closure test for \mathbb{L}_T

$x = a^{k-1}$	$wxy = \pi a^{k-1} u$	$\in \mathbb{L}_T$
	$vxz = \varepsilon a^{k-1} \varsigma$	$\in \mathbb{L}_T$
	$wxz = \pi a^{k-1} \varsigma$	$\in \mathbb{L}_T$
	$vxy = \varepsilon a^{k-1} u$	$\notin \mathbb{L}_T$

The Tutrugbu blocking pattern also cannot be SP_k for any k because it violates Subsequence Closure (SC).

Table 15: Subsequence Closure test for \mathbb{L}_T

$w = \varepsilon \varepsilon a^{k-1} u o$	$\in \mathbb{L}_T$
$v = \varepsilon u \sqsubseteq w$	$\notin \mathbb{L}_T$

While Table 14 showed that the \mathbb{L}_T blocking pattern is necessarily a long distance pattern, the test in Table 15 illustrates why circumambience is more complex than Finnish transparency, for example. The Tutrugbu [+low] vowel does block [+ATR] harmony regardless of how far either the blocker or a [-low] prefix vowel is from the [+ATR] vowel. However, two triggers are needed for a [+low] vowel to block harmony: a [+ATR] vowel on the right of the blocker and a [+high] vowel at the left edge of the string. The example test in (15) shows that requiring the [+high] vowel to be string-initial allows a grammatical string to also contain a subsequence which is ungrammatical on its own.

The UCA pattern in Tutrugbu is also not TSL_k for any k because it violates SSC even over a tier projection. The definition of TSL states that a tier only contains a strict subset of the elements on the original string, i.e. vowel string (Heinz, Rawal, and Tanner 2011). The alphabet for Tutrugbu vowel strings contains all the Tutrugbu vowels $\Sigma = \{i, e, o, \iota, \upsilon, \varepsilon, \varsigma, a\}$ and the TSL tier would contain all the [+high], [+low], and [+ATR] vowels $T = \{i, e, o, \iota, \upsilon, a\}$. Over T the Tutrugbu pattern is not SL_k for any k because it does not obey SSC, as shown in Table 16.

Table 16: Suffix Substitution Closure test over T for \mathbb{L}_T

$x = a^{k-1}$	$wxy = a^{k-1} \upsilon$	$\in \mathbb{L}_T$
	$vxz = \iota a^{k-1} u$	$\in \mathbb{L}_T$
	$wxz = a^{k-1} u$	$\notin \mathbb{L}_T$

Table 16 illustrates that even without the [-ATR, -high] vowels on T , \mathbb{L}_T contains vowel strings with full [-ATR] harmony like in wxy if w is the empty string and vowel strings with UCA blocking like in vxz . However, when the suffixes are switched between wxy and vxz the result is the disharmonic string with a [-ATR, +low] vowel succeeded by a [+ATR, +high] vowel in wxz . The resulting string wxz is ungrammatical in Tutrugbu because there is no initial [+high] vowel to create the environment for blocking which would justify the surface disharmony.

The UCA Tutrugbu blocking pattern is locally testable (LT) over strings of vowels because it obeys Local

Test Invariance (LTI) and is describable using propositional logic enriched with a string boundary symbol. While CNLs consist of only a set of restrictions or negative literals about what is logically possible in a grammar, propositional logic adds logical connectives such as the conditional and can be used to connect entire statements. The propositional logic statement that describes the Tutrugbu blocking pattern utilizes an additional string boundary symbol and a conditional to more precisely describe the environment in which blocking cannot occur. That statement is then conjoined to a CNL which describes the basic harmony pattern, as shown below in (100).

(100) Propositional logic statement of Tutrugbu blocking pattern

$$\begin{aligned} & (\neg[+ATR,-low][\neg ATR,-low] \wedge \neg[\neg ATR,-low][+ATR,-low]) \wedge ((\#[\neg high]) \Rightarrow (\neg[+low][+ATR])) = \\ & (\neg i\bar{i} \wedge \neg i\bar{u} \wedge \neg i\bar{\epsilon} \wedge \neg i\bar{o} \wedge \neg i\bar{a} \wedge \neg u\bar{i} \wedge \neg u\bar{u} \wedge \neg u\bar{\epsilon} \wedge \neg u\bar{o} \wedge \neg u\bar{a} \wedge \neg \epsilon\bar{i} \wedge \neg \epsilon\bar{u} \wedge \neg \epsilon\bar{\epsilon} \wedge \neg \epsilon\bar{o} \wedge \neg \epsilon\bar{a} \wedge \neg o\bar{i} \\ & \wedge \neg o\bar{u} \wedge \neg o\bar{\epsilon} \wedge \neg o\bar{o} \wedge \neg o\bar{a} \wedge \neg i\bar{i} \wedge \neg u\bar{i} \wedge \neg \epsilon\bar{i} \wedge \neg o\bar{i} \wedge \neg u\bar{u} \wedge \neg \epsilon\bar{u} \wedge \neg o\bar{u} \wedge \neg i\bar{\epsilon} \wedge \neg u\bar{\epsilon} \wedge \neg \epsilon\bar{\epsilon} \wedge \neg o\bar{\epsilon} \\ & \wedge \neg i\bar{o} \wedge \neg u\bar{o} \wedge \neg \epsilon\bar{o} \wedge \neg o\bar{o}) \wedge ((\#e \vee \#o \vee \#\epsilon \vee \#o \vee \#a) \Rightarrow (\neg ai \wedge \neg au \wedge \neg ae \wedge \neg ao)) \end{aligned}$$

The logical statement in (100) is read as “A string must not contain any of the two-factors in which two [-low] vowels have different ATR feature values AND if a string begins with one of the [-high] vowels [e, o, ε, o, or a] then it must not contain any of the two-factors in which [+low] [a] is followed by a [+ATR] vowel, namely [ai, au, ae, or ao]”. The featural CNL is simply a shorthand to help clarify full vocalic CNL. The statement in (100) combines the conditional with a SL₂ CNL in order to describe both the general harmony pattern and the UCA blocking pattern. The conditional within this statement succinctly describes precisely the environment in which blocking can not occur, which allows strings with UCA blocking and a potentially unbounded number of blockers. The combination of two statements within the conditional describes the circumambient environment of blocking without referring to blocking directly because a string cannot contain an initial [-high] vowel and disharmony with [a]. In other words blocking—[a] followed by a [+ATR] vowel—is only possible when the initial vowel is also [+high].

As an input-output mapping process Tutrugbu has been analyzed as unique and more complex than most other phonological patterns (Chandlee 2014; Heinz and Lai 2013; McCollum et al. 2020) but in this section I have shown that on the surface Tutrugbu is more like other phonological patterns than previously thought. Over multi-tiered ARs Tutrugbu ATR harmony with UCA blocking is diffuse just like Akan and describable with a set of four FSCs over multi-tiered ARs which means that it fits into the ASL^{VH} class of ARsets. Over strings of vowels Tutrugbu ATR harmony with UCA blocking falls into the LT class of stringsets which makes it more complex than the SL_k and SP_k phonotactic patterns discussed in the previous chapter. So as a phonotactic pattern Tutrugbu is not so unique and over multi-tiered ARs it fits into my new class along with many other vowel harmony patterns. These findings demonstrate how an enriched representation like the multi-tiered ARs

used here reduces the computational complexity of a pattern. UCA blocking is at least LT and its description requires propositional logic over strings of vowels but the same pattern is local and describable with a CNL over multi-tiered ARs.

5.2 Eastern Meadow Mari

Eastern Meadow Mari (EMM) is a Uralic language spoken in Russia. This language provides a counter-example to my ASL hypothesis and highlights a difference between the computations of blocking and transparent neutral vowels. The EMM vowel harmony pattern is more complex than other patterns analyzed in this dissertation because it utilizes transparent vowels and those vowels contrast in the harmonic feature. In addition, EMM exhibits positional transparency which means these two vowels only behave as transparent when they occur in non-initial positions of a word.

EMM utilizes a 9-vowel system with distinctions in height, rounding, and backness. This section focuses on backness harmony; so the [-back] vowels are [i, y, e, ø, æ] and the [+back] vowels are [u, o, ə, a]. The [+high] vowels are [i, y, u], the [-high] vowels are [e, ø, ə, o, æ, a], the [+round] vowels are [y, ø, u, o], and the [-round] vowels are [i, e, æ, ə, a] (Vaysman 2009; Walker 2011).

Table 17: Eastern Meadow Mari Vowels

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

The positionally transparent vowels are [e, ə] and these will be discussed in more detail later.

EMM uses backness harmony and the two [-high, -low, -round] vowels [e, ə] behave as transparent only in non-initial positions. So EMM words without a medial or final [e] or [ə] exhibit full backness harmony, as in (101). All of the EMM data presented in this dissertation is taken from Vaysman (2009).

(101) Eastern Meadow Mari full back harmony (Vaysman 2009)

a. Roots

i. *fyzær* ‘sister’

ii. *murna* ‘tube, pipe’

b. Nom.sg.2p.pl.poss

i. *em†dæ* ‘your (pl) medicine’

ii. *tʃødræ†tæ* ‘your (pl) forest’

iii. *kutko†ta* ‘your (pl) ant’

c. Nom.sg.3p.pl.poss

- i. ij+næ ‘our year’
- ii. tynæ+næ ‘our world’
- iii. tam+na ‘our taste’

d. Dative

- i. jør+læn ‘milk (dat.)’
- ii. lum+læn ‘snow (dat.)’
- iii. kawun+læn ‘pumpkin (dat.)’

Non-initial [e] and [ə] are transparent so vowels on either side of them must have the same back feature value. In word-final position these two vowels occur interchangeably with [+back] and [-back] words. Crucially, it is in word-initial position where these vowels’ behavior changes. For example the suffix vowel in (101bi) is [-back] just like the initial [e]. In (102ai-ii, iv) and (102bii-iv) we see that the first and second vowels are both [-high, -low, -round] but they have different back feature values.

(102) Eastern Meadow Mari transparency (Vaysman 2009)

a. Roots

- i. teŋgəz ‘sea’
- ii. serəf ‘letter’
- iii. pəreŋgə ‘tree’
- iv. jəlme ‘tongue, language’

b. Dative

- i. imjə+læn ‘horse (dat.)’
- ii. serəf+læn ‘letter (dat.)’
- iii. təlze+læn ‘moon, month (dat.)’
- iv. ləβe+læn ‘butterfly (dat.)’

The suffixes in (102bii-iv) have the same back feature value as the initial [-high, -low, -round] vowel and so harmony appears to have skipped over the medial [-high, -low, -round] vowel.

The [-high, -low, -round] vowels have both a [-back] and a [+back] counterpart, but they could still be considered neutral according to the definition in chapter 4. I call these vowels positionally transparent because they require harmony on either side of them when they occur in word-medial positions and they can occur interchangeably at the end of words with harmony in either back feature value.

The two sources for EMM disagree on some aspects of the pattern. Vaysman (2009) claims that EMM uses full back harmony which is affected by full/reduced vowel alternations and metrical structure. On the other hand, Walker (2011) describes the EMM back harmony pattern as looking like first-last harmony (FLH) because for certain suffixes “vowels in the word-final syllable assimilate in backness to the vowel in the initial syllable”, but “vowels that intervene between the initial and final one can be transparent.” Previous work claimed that FLH is unattested because it is computationally too complex; a true FLH description would require at least propositional logic over strings (Heinz 2018; Lai 2015).⁷

⁷C’Lela includes a height harmony pattern which has also been cited as a potential example of FLH but Lai (2015) points out that the data is consistent with a variety of interpretations. In addition, the height harmony appears to apply only with one circumfix and over a bounded distance. This pattern thus warrants further investigation but does not present a clear attestation of FLH.

However, the computational perspective and meticulous methods I employ throughout this dissertation proved especially useful to clarify the EMM back harmony pattern. I first listed out all the data provided in both Walker (2011) and Vaysman (2009), then I listed only the vowel strings found in roots and in words with any suffixes. While the two descriptions of EMM back harmony appear to contradict each other, close inspection of the surface forms and vowel strings attested in EMM revealed a pattern that actually combines the two descriptions: EMM does utilize full back harmony and two specific vowels are transparent only when they are in a non-initial position. Walker (2011) was on the right track including EMM in a larger discussion of positional licensing because while it does not exhibit FLH as she suggested, the transparency of [e] and [ə] is subject to a positional restriction. In short the vowel harmony pattern found in EMM surface forms can be described by the following insights:

- EMM words consist of roots with up to three syllables
 - followed by one or two suffixes with up to two syllables.
- Two vowels [e] and [ə] are transparent in non-initial positions
 - non-alternating suffixes all contain only [e] or [ə]
 - also occur at the beginning of a word followed only by vowels with the same back feature value

The second point is what makes EMM back harmony unique. The two [-high, -low, -round] vowels [e] and [ə] each occur between both [-back] and [+back] vowels in the same way that transparent vowels like those in Finnish can occur between vowels with any back feature value.⁸ Another piece of evidence for the transparency of [-high, -low, -round] vowels [e] and [ə] is that these are the only two vowels that occur in non-alternating suffixes and thus each follows both roots with [-back] vowels and roots with [+back] vowels. In addition, the [-high, -low, -round] vowels do not precede vowels of both backness specifications when they occur at the beginning of words which makes it look like they are no longer transparent in initial position.

The non-ASL^{VH} complexity of EMM illustrates the effect that neutral vowel contrasts have on vowel harmony patterns with transparency. Vowel harmony patterns with a contrast in the harmonic feature of transparent vowels are necessarily more complex than transparent vowels with no harmonic contrast. For example, in Finnish back harmony is local because a single [-back] feature is diffused across any span of transparent vowels no matter how long it is. However, the contrast in backness between EMM's transparent vowels means that a back feature iterates across a span of alternating transparent vowels. This means that the distance between harmonizing back features depends upon the number of alternating transparent vowels


⁸There is one root which appears to suggest that the [+low] vowel [a] can also be transparent: [merəŋ] 'hare', [merəŋ+læn] 'hare (dative)', [merəŋ+ge] 'hare (comitative)' (Vaysman, 2009; p70). However, these examples are contradicted by a different transcription of the same root [meræŋ] 'hare' (Vaysman, 2009; p80). In addition, this is the only such example of a word with [e] or any other [-back] vowel followed by [a]. I thus assume the grammatical form of this root is [meræŋ] and [ea] is not a grammatical string of vowels in EMM.

used. This dependency makes transparency with a harmonic contrast more complex than transparency with no harmonic contrast and the reasoning is presented in more detail below. In the next subsection I present my analysis of EMM back harmony with positional transparency over multi-tiered ARs and show why this pattern is not ASL^{VH} . The following subsection presents my analysis of EMM back harmony with positional transparency over strings of vowels.

5.2.1 EMM transparency over multi-tiered ARs

EMM back harmony involves transparency with a positional restriction and in this section I show that the EMM pattern does not fit into the ASL^{VH} class of ARsets as I have defined it. The EMM pattern provides evidence of a difference in computational complexity between strings and ARs as well as between blocking and transparency patterns in vowel harmony. While Tutrugbu combined basic blocking FSCs with some additional FSCs over multi-tiered ARs to capture the positional restriction the EMM back harmony pattern cannot be described so simply. The combination of iterativity across an unbounded distance, requiring reference to a position that succeeds a word boundary, and the possibility of a vowel string containing multiple alternating transparent vowels prevents the positional transparency pattern from being describable with a set of connected FSCs over multi-tiered ARs that use only one ordering relation. Crucially, the back contrast amongst transparent vowels predicts that a vowel string could contain a span of transparent vowels with different back feature values, as in (103).

(103) Possible non-local EMM AR




While this AR clearly shows iterative locality on the round tier it is the back tier which defines the locality of the back harmony pattern. So unlike in Finnish, a single back feature is not multiply associated to all the transparent vowels between two harmonizing vowels. Rather, each transparent vowel is associated to a different back feature which makes the EMM back harmony pattern non-local. This combination of facts about EMM positional transparency prevents the pattern from fitting into the ASL^{VH} class.

Jardine (2019) defines ASL patterns as being forbidden substructure grammars (FSG) over ARs. First he proves in Lemma 1 that for each FSG there is an equivalent statement in first order logic (FO). FO is the level of logic which describes exactly the Star Free (SF) stringsets. Jardine (2019) then goes on to prove that the ASL class is a strict subset of the SF class because all ASL patterns are also describable as SF patterns, but not

as FLH. The results of a phonetic or a computational investigation would benefit phonological theory as a whole but unfortunately these are beyond the scope of this dissertation.


Regardless, EMM already provides valuable insights for phonological theory because it highlights the importance of neutral vowel contrasts for the complexity of patterns with transparency. You may recall from chapter 4 that Finnish exhibits back harmony with transparent vowels that do not contrast in backness: [i] and [e]. As transparent vowels these two can occur interchangeably in a sequence; and regardless of how many occur within a single span all the transparent vowels are associated to a single [-back] feature, as shown in (107) below. The multiple association of transparent vowels to a single [-back] feature makes Finnish back harmony local over multi-tiered ARs and this locality is the reason that this pattern fits into the ASL^{VH} class.

(107) Finnish: [maisemia] ‘scenery.plural.partitive’



On the other hand, the contrast in backness between the two positionally transparent EMM vowels [e] and [ə] prevents such locality. The two vowels can still occur interchangeably because they are considered to be transparent but each vowel is associated to a different back feature as shown in (103) above and repeated in (108) for clarity. Because a span of positionally transparent EMM vowels is not multiply associated to a single back feature adding more alternating transparent vowels to a span would also increase the number of back features present in the AR.

(108) Possible non-local EMM AR



So unlike Finnish the EMM back harmony pattern is not local over multi-tiered ARs because its positionally transparent vowels contrast in backness. This non-locality is what prevents EMM from fitting into the ASL^{VH} class. In short, EMM shows that in order to claim all vowel harmony is ASL^{VH} transparent vowels cannot contrast in the harmonic feature.

5.2.3 EMM over strings of vowels

The set of strings of vowels that represents the back harmony pattern in EMM is represented as \mathbb{L}_{EMM} . Examples of some of the logically possible strings that occur in \mathbb{L}_{EMM} and those that do not are shown in Table 18. The back harmony pattern and the positional transparency of [e] and [ə] apply both within roots and in words with one or more suffix(es) so it is not necessary to represent the morpheme boundaries in strings of \mathbb{L}_{EMM} .

Table 18

	strings $\in \mathbb{L}_{EMM}$	strings $\notin \mathbb{L}_{EMM}$
Full Harmony:	eø, yæ, ua, əo iæø, øeæ, uoa, aəo	ea, øa, oæ, aø eyo, iøa, uoæ, oəø
Transparency:	iə, yəæ, ae, ueo	iəo, øəa, oəæ, ueø

The EMM back harmony pattern with positional transparency cannot be SL_k for any k because it violates SSC.

Table 19: Suffix Substitution Closure test for \mathbb{L}_{EMM}

$x = ə^{k-1}$	$wxy = iə^{k-1}ø$	$\in \mathbb{L}_{EMM}$
	$vxz = oə^{k-1}a$	$\in \mathbb{L}_{EMM}$
	$wxz = iə^{k-1}a$	$\notin \mathbb{L}_{EMM}$

The test in Table 19 illustrates that the restriction on backness of harmonizing vowels holds regardless of how many transparent vowels might intervene. Swapping the suffixes of two strings in \mathbb{L}_{EMM} that involve harmony in different back feature values results in two strings that are not in \mathbb{L}_{EMM} because the harmonic back feature no longer iterates across the transparent vowel(s). Thus back harmony holds across any arbitrary distance so it cannot be SL_k for any k .

EMM back harmony with positional transparency also cannot be SP_k for any k because it violates SC.

Table 20: Subsequence Closure test for \mathbb{L}_{EMM}

$w = əuae^{k-1}o$	$\in \mathbb{L}_{EMM}$
$v = eo \sqsubseteq w$	$\notin \mathbb{L}_{EMM}$

The test in Table 20 shows the computational effect of neutral vowels with a harmonic contrast on a harmony pattern with transparency. For example, in Finnish the transparent vowels had no harmonic counterpart so any (sub)string that began with [i] or [e] was grammatical regardless of the back feature value of the successive vowels. However, EMM has transparent vowels with a harmonic contrast and they participate in harmony in initial position; so a (sub)string which begins with [e] or [ə] is ungrammatical if the successive harmonizing vowel(s) do not have the same back feature value. In other words, a back feature can iterate across non-initial

[e] and [ə], but in initial position these vowels are not transparent.

EMM back harmony with positional transparency is also not TSL_k for any k because a TSL tier cannot distinguish between vowels in different positions of a vowel string. The alphabet for EMM contains all the EMM vowels $\Sigma = \{i, y, e, \emptyset, \text{æ}, u, \text{ə}, o, a\}$ and a tier which excludes the non-initial transparent vowels [e, ə] would also necessarily exclude the harmonizing initial [-high, -low, -round] vowels: $T = \{i, y, \emptyset, \text{æ}, u, o, a\}$. By removing the possibility of writing constraints which enforce harmony with initial [e] and [ə] a SL CNL for EMM written over T would necessarily allow ungrammatical disharmonic strings such as *eoua* or *əyøæ*. Treating [e] and [ə] as unconditionally transparent by removing them from a TSL tier is thus the wrong approach for EMM. A complete description of positional transparency requires the inclusion of constraints on the possible vowels which can succeed an initial [-high, -low, -round] vowel in order to enforce harmony. TSL tiers cannot refer to positions within a string which makes it impossible to describe EMM in a local way over T .

The importance of the initial position for the otherwise transparent vowels [e] and [ə] in EMM causes this back harmony pattern to be LT over strings of vowels because its description requires propositional logic enriched with word boundary symbols.

(109) Propositional logic statement of EMM back harmony with transparency

$$\begin{aligned}
 & (\#[\text{-back, -high, -low, -round}] \Rightarrow \neg([\text{+back, +round}] \vee [\text{+back, +low}])) \wedge \\
 & (\#[\text{+back, -high, -low, -round}] \Rightarrow \neg([\text{-back, +high}] \vee [\text{-back, +round}] \vee [\text{-back, +low}])) \wedge \\
 & \neg([\text{+back, +high}] \wedge [\text{-back, +high}]) \vee ([\text{+back, +round}] \wedge [\text{-back, +round}]) \vee \\
 & ([\text{+back, +round}] \wedge [\text{-back, +low}]) \vee ([\text{+back, +low}] \wedge [\text{-back, +high}]) \vee \\
 & ([\text{+back, +low}] \wedge [\text{-back, +round}]) \vee ([\text{+back, +low}] \wedge [\text{-back, +low}]) \\
 & = (\#e \Rightarrow \neg(u \vee o \vee a)) \wedge (\#\text{ə} \Rightarrow \neg(i \vee y \vee \emptyset \vee \text{æ})) \wedge \\
 & \neg((u \wedge i) \vee (u \wedge y) \vee (u \wedge \emptyset) \vee (u \wedge \text{æ}) \vee (o \wedge i) \vee (o \wedge y) \vee (o \wedge \emptyset) \vee (o \wedge \text{æ}) \vee (a \wedge i) \vee (a \wedge y) \vee (a \wedge \emptyset) \vee (a \wedge \text{æ}))
 \end{aligned}$$

The statement in (109) reads colloquially as the following: “If a string begins with [e] then it cannot contain any of [u], [o], or [a] and if a string begins with [ə] then it cannot contain any of [i], [y], [ø], or [æ]. In addition, no string can contain both a [+back] and a [-back] vowel, unless it is non-initial [e] or [ə].” The featural statement at the top is intended solely to be a shorthand which clarifies the vocalic statement below it. However, just as with transparency in Finnish the need to refer to multiple features makes the featural statement longer and less clear than the vocalic one. The first portions of each statement are interpreted with the successor relation (\triangleleft) and utilize the conditional to clarify which vowels can cooccur in a string with an [e] that succeeds a word boundary and in strings with a [ə] that succeeds a word boundary. These are conjoined with a more general statement forbidding certain vowel combinations in a string, which reinforces

the prevalence of back harmony throughout a word, even across transparent vowels. The last portion utilizes negation and does not require interpretation with either a successor or precedence relation, but it accomplishes a goal similar to that of a CNL interpreted with precedence.⁹

In this section I have shown that EMM back harmony utilizes positional transparency with contrast in backness between transparent vowels and all of this makes EMM back harmony more complex than other vowel harmony patterns over both multi-tiered ARs and strings of vowels. The EMM pattern provides a counterexample to the hypothesis that all vowel harmony patterns are ASL^{VH} . In order to justify maintaining the ASL^{VH} hypothesis a restrictive theory of vowel harmony would have to define transparency as excluding vowels with a harmonic contrast. This more restrictive definition of transparency would also align with the more conservative definition of neutral vowels discussed in the previous chapter.

5.3 Conclusion

In this chapter I have shown that vowel harmony patterns which previous work considered non-local and more complex than SL or SP reveal important distinctions between the computations of different types of neutral vowels. First, Tutrugbu exhibits UCA blocking, which can be described locally over multi-tiered ARs and EMM exhibits positional transparency which cannot. I argue that different types of neutral vowels—blocking vs. transparent—cause vowel harmony patterns to fall into different complexity classes. In short, transparency is more complex than blocking. In chapter 4 I showed that over strings the basic blocking pattern in Akan is SL and the basic transparency pattern in Finnish is SP. In this chapter I have shown that even over multi-tiered ARs, EMM back harmony with positional transparency is more complex than Tutrugbu ATR harmony with UCA blocking. My claim that there is a difference in complexity between blocking and transparency is further supported by the fact that the description of transparent vowels in both Finnish and EMM necessarily refers to more than one feature. This means the NLs over strings of features refer to sets of feature combinations and so the featural CNLs for vowel harmony patterns with transparency in Finnish and EMM are much longer than the vocalic CNLs for the same patterns; whereas the featural CNLs for vowel harmony patterns with blocking in Akan and Tutrugbu were much shorter and more clear than their vocalic counterparts. While the featural CNLs were intended solely to be a shorthand to help clarify the pattern found in vowel strings they turned out not to be shorter or more clear for patterns with transparency.

Second, EMM positional transparency is more complex than other vowel harmony patterns with transparent vowels like Finnish because EMM transparent vowels contrast in the harmonic feature. EMM has two

⁹I was able to write a CNL that encompasses the restrictions in the first two conditional statements as well as one for the final statement, but the first CNL was interpreted with successor and the second was interpreted with precedence. The overall EMM back harmony pattern requires propositional logic in order to combine general back harmony and transparency with the positional restriction on the otherwise transparent vowels [e] and [ə].

transparent vowels [-back] [e] and [+back] [ə]. My theory predicts that a span of successive transparent vowels could include both [e] and [ə]; and such a span would not be diffuse like it would in Finnish because all the transparent vowels would be associated to different back features with alternating values. Thus even on the back tier an EMM AR could be non-local. So while both Tutrugbu and EMM are relatively complex over strings their analyses over multi-tiered ARs provide crucial evidence of the computational effects that neutral vowels have on vowel harmony patterns in general; vowel harmony patterns that include transparent vowels with a harmonic contrast are more complex than vowel harmony patterns that use transparent vowels with no harmonic contrast. In short, if one claims that all vowel harmony is ASL^{VH} then transparent vowels cannot contrast in the harmonic feature.

6 ARs provide concise descriptions of patterns

This chapter provides evidence for the second argument in favor of using multi-tiered ARs to represent vowel harmony patterns: they allow the descriptions of patterns to be more concise. I present two vowel harmony patterns which are SL but require extensive CNLs over strings. Both patterns can be described with far fewer restrictions by using multi-tiered ARs. Baiyinna Orochen presents a surface rounding harmony pattern which relies on vowel height, length, and blocking. This pattern is describable with a SL_3 CNL that must contain at least 56 cooccurrence restrictions over strings of vowels. However, Baiyinna Orochen is describable with a set of only three FSCs over multi-tiered ARs. Kinande presents a surface vowel harmony pattern which was previously described as having one pattern with some exceptional words. I show that all the Kinande data can be described by a single pattern using a SL_2 CNL with 33 cooccurrence restrictions over strings of vowels or a CNL of five FSCs over multi-tiered ARs. While both vowel harmony patterns presented in this chapter are SL over strings of vowels they can be described using far fewer restrictions when they are represented over multi-tiered ARs.

6.1 Baiyinna Orochen

Baiyinna Orochen is a Tungusic language “spoken in the village of Baiyinna, Huma County, Heilongjiang Province, in northeastern China”(Li 1996, 6; Walker 2014a, 503). B. Orochen utilizes sixteen monophthongs and two diphthongs. These eighteen vowels are distinguished by height, backness, rounding, and tongue root position; and the monophthongs exhibit a length distinction. The [+round] vowels are also all [+back] [u, u:, ʊ, ʊ:, o, o:, ɔ, ɔ:], but there are both [-round, -back] vowels [i, i:, ɪ, ɪ:, ie, iɛ] and [-round, +back] vowels [ə, ə:, a, a:] in B. Orochen.

Table 21: Baiyinna Orochen Vowels

	-back	+back		
+high	i, i:		u, u:	+ATR
	ɪ, ɪ:		ʊ, ʊ:	-ATR
-high	\widehat{ie}	ə, ə:	o, o:	+ATR
	$\widehat{iɛ}$	a, a:	ɔ, ɔ:	-ATR
	-round		+round	

B. Orochen utilizes both ATR and rounding harmony, but Walker (2014a) analyzes only the round harmony pattern so I will also focus only on the rounding harmony in B. Orochen. The features relevant for the round harmony pattern in B. Orochen are [\pm round], [\pm high], and [\pm long] so there are six crucial featural distinctions for this harmony pattern: [-round, +high, -long] [i, ɪ]; [-round, +high, +long] [i:, ɪ:]; [-round, -high, -long] [ie, iɛ, ə, a]; [-round, -high, +long] [ə:, a:]; [+round, +high, -long] [u, ʊ]; [+round, +high, +long] [u:, ʊ:];

[+round, -high, -long] [o, ɔ]; and [+round, -high, +long] [o:, ɔ:].

The surface round harmony pattern in B. Orochen can be described as follows: a word that begins with an [o] will contain only [o] or [o:] vowels thereafter, a word that begins with [ɔ] will contain only [ɔ] or [ɔ:] vowels after, and [+high] vowels block round harmony. B. Orochen blockers exhibit a contrast in the harmonic feature because the vowel inventory above includes both [+high, -round] [i, i:, ɪ, ɪ:] and [+high, +round] [u, u:, ʊ, ʊ:]. I have argued in previous chapters that such a contrast does not preclude naming blocking vowels as neutral and that these neutral vowel contrasts do not increase the complexity of vowel harmony patterns with blocking, but they can do so with transparency. Some examples with full diffuse [+round] harmony are shown in (110) below. All of the B. Orochen data presented in this chapter comes from Li (1996).

(110) Round harmony with [-high] root vowels (Li 1996)

- | | |
|----------------------------|---------------------------------|
| a. tʃoloɒon ‘morning star’ | e. sɔbɒɔ ‘fish skin’ |
| b. moyon ‘silver’ | f. ɔɔktɔ ‘hay’ |
| c. sokko: ‘muddy(water)’ | g. ɒɔɔ ‘log’ |
| d. olo:k ‘lie’ | h. ɔmɔ:ŋ ‘fatty meat (of deer)’ |

This [+round] harmony holds both within roots and across morpheme boundaries from roots to suffixes, as the words in (111) show.

(111) Round harmony with [-high] suffix vowels (Li 1996)

- | | |
|--|---------------------------------|
| a. sokko: + mɒɔ ‘muddy(water) (contem)’ | e. ɔm + ma ‘who likes to drink’ |
| b. sosok + jɔ ‘pasture (indef.acc)’ | f. ɔɔ + jɔ ‘mountain’ |
| c. ɔmɔ:ŋ + mɔ ‘fatty meat of deer (def.acc)’ | g. ɔwi + mɔ ‘who likes to play’ |
| d. ɔɔ + jɔ ‘fish (indef.acc)’ | h. bɪɔ + ma ‘river’ |

The words in (111a-d) all exhibit [+round] harmony because the suffixes include a [+round] vowel following a root with only [+round] vowels. The words in (111e-h) have the same suffix but following a [-round] root the suffix vowel surfaces as [-round]. You may also notice that in (111e-f) the suffix vowel surfaces as [-round] following the [-round, -high] root vowel despite containing an iinitial [+round, +high] vowel. These [+high] neutral vowels will be discussed in more detail later but what matters in these examples is that the suffix vowel participates in rounding harmony and surfaces with a [±round] vowel depending upon the round feature value of the root vowel(s).

An analysis of the round harmony pattern in B. Orochen must also account for words that begin with any vowel other than [o] or [ɔ]. The examples in (112) illustrate full [-round] harmony.

(112) Words with initial [-round] vowels (Li 1996)

- | | |
|---------------------------|--|
| a. əgdəgə ‘big’ | e. bandən+mə ‘bench (borrowed from Chinese)’ |
| b. targan ‘field, garden’ | f. fafə+lə:tʃə ‘to launch’ |
| c. sə:ksə ‘blood’ | g. ʃəntʃan+lə:tʃa ‘to manufacture’ |
| d. tʃa:lban ‘birch’ | |

These examples all begin with [-round, -high] vowels and contain only [-round, -high] vowels. Words with [±round, +high] vowels will be discussed later with respect to blocking.

6.1.1 [+long] vowels

The round harmony pattern in B. Orochen is relatively unique because vowel length is a factor in whether or not the [+round] feature is diffused. On the surface [+round, -high] vowels only occur in non-initial positions if the initial vowel is [+round]; and in fact the initial vowel must also be [-high, -long]. The representation of phonological vowel length often differs from other features because it so rarely interacts with segmental patterns like vowel harmony. In order to explain tone, vowel shortening, and compensatory lengthening patterns a variety of arguments have been made over the years for treating phonological vowel length as either a sequence of identical vowels (Kenstowicz 1970; McCawley 1968; Odden 1987, 1996) or as a SPE (Chomsky and Halle 1968) style feature of the vowel (Fidelholtz 1971; Kenstowicz 1970; Pyle 1971). For B. Orochen specifically, Li (1996) transcribes phonological vowel length as two consecutive identical vowels. This representation removes vowel length from the set of subsegmental features that can be referenced in constraints. Representing length as the duplication of a short vowel prevents a theory from being able to distinguish short and long vowels using phonological features.


Autosegmental theories have also been proposed which account for vowel length as a binary prosodic category (Clements 1986; Clements and Keyser 1983; Ingria 1980; Keyser and Kiparsky 1984; Levin 1985; McCarthy 1979, 1982; White 1972) and in the late 1980s-1990s phonologists began to explore moraic theories of vowel length in order to explain syllable weight in stress patterns (Hayes 1989; Hyman 1985; McCarthy and Prince 1986; Zec 1988). Following a comprehensive summary of this theoretical history, Odden (2011) concludes that phonological length must be distinguished from other vowel properties and this is best done with an autosegmental theory which associates one thing (i.e. segment) to two higher-level things (i.e. moras). For B. Orochen, while Walker (2014a) does use moras she explicitly encodes vowel length as a subscripted

mora on the vowel, such as in her constraint \forall -Harmony([round]/V_μ[-high], V) (p. 512). In this way Walker (2014a) is explicitly treating vowel length as part of the makeup of the vowel much like other vowel features such as [±high] and [±round].


Unlike most of the evidence cited in Odden (2011) for the prosodic account of vowel length, B. Orochen clearly uses phonological vowel length to interact with other vowel features in what is considered a segmental pattern: rounding harmony. In addition, my theory must be able to clearly distinguish the vowel strings of words with short vowels in multiple consecutive syllables (ooo) from those with long vowels (o:o) and using Li (1996)'s representation makes this distinction impossible. I thus adopt a representation of vowel length based on Walker (2014a) by treating it as an explicit part of the makeup of a vowel which I represent as a feature [±long] that occupies its own autosegmental tier and is directly associated to each vowel, as shown in (113).

(113) [kɔ:ŋakta] 'handbell' (Li 1996)

a. Walker (2014)'s representation



b. My representation



The AR in (113a) shows the representation of vowel length that Walker (2014a) uses with subscripted moras to distinguish long from short vowels. The AR in (113b) illustrates how I represent vowel length as another binary vowel feature, which is directly associated to the vowel(s). The length feature is a direct representation of the IPA length marker [ː] and using a length feature allows me to describe the B. Orochen pattern with only three FSCs over multi-tiered ARs. More work on the interaction of vowel length with vowel harmony patterns is needed to discern the precise effects of different vowel length representations.

In B. Orochen specifically the two [+long, +round, -high] vowels [o: and ɔ:] exhibit a unique behavior on the surface. The descriptions in Walker (2014a) and Li (1996) state that they do not occur initially in words with only [+round] vowels. However, they can occur in non-initial positions within words containing [+round] vowels and in initial position within words containing [-round] vowels. These two vowels thus appear to only participate in rounding harmony when they are in non-initial positions of a word. Some examples of B. Orochen words with initial [+long,+round,-high] vowels are shown in (114) below. Unlike in (110)-(111), the words in (114) all begin with [+round, +long] vowels succeeded by [-round] vowels.

(114) Words with initial [+long,+round,-high] vowels (Li 1996)

- | | |
|---|---------------------------------|
| a. o:ɗən ‘velvet’ | e. bo:l+ɟə ‘slave (indef.acc)’ |
| b. ko:məxə ‘windpipe’ | f. bo:l+wə ‘slave (def.acc)’ |
| c. kə:ŋakta ‘handbell’ | g. go:l+ja ‘policy (indef.acc)’ |
| d. tɔ:lga ‘pole used to support coffin’ | h. kɔ:+xɑ:n ‘wine pot (dim)’ |

As described, all of the words in (114) begin with [o:] or [ɔ:] but all the vowels after are [-round]. In addition the non-initial behavior of these two vowels is shown in (115).

(115) Words with non-initial [+round,+long,-high] vowels (Li 1996)

- | | |
|---------------------------------|---|
| a. ɔpɔ: ‘rocky hillock’ | h. gɔlɔ:+tkɔ:xi ‘log (direct)’ |
| b. ɔrɔ:n ‘top, surface’ | i. ɔpɔ:+lɔ: ‘rocky hillock (destin)’ |
| c. gɔlɔ: ‘log’ | j. ɔmɔ:ŋ+mɔ ‘fatty meat of deer or roe deer (def.acc.)’ |
| d. ɔmɔ:ŋ ‘fatty meat (of deer)’ | |
| e. olo:- ‘to cook’ | k. olo:+no+tʃo ‘to cook (intent.asp.pt.t) |
| f. olo:k ‘lie’ | l. sokko:+mɲo ‘muddy(water) (contem)’ |
| g. sokko: ‘muddy (water)’ | m. dokto+lɔ:+ro ‘to harness (prt)’ |

Whether word-medially or word-finally [o:] and [ɔ:] are surrounded by [+round, -high] vowels. So even though the initial vowel is [+round, -high] it must be short if it is followed by other [+round, -high] vowels, including [+long] ones.

B. Orochen exhibits basic rounding harmony amongst [-long, +round, -high] vowels on the surface and also includes [+long, +round, -high] vowels with some unique (dis)harmonic behavior. In addition, B. Orochen includes [+high] vowels and these are discussed in the following subsection.

6.1.2 [+high] vowels


B. Orochen also has neutral vowels which are [+high]. These neutral [+high] vowels are unique in that they exhibit both blocking and transparency on the surface. Within words containing [+round, -high] vowels the [+high] vowels block [+round] harmony. Within words containing [-round, -high] vowels the [+high] vowels are transparent to [-round] harmony. So in B. Orochen [+round] is diffused up to and [-round] is iterated across [+high] vowels.

B. Orochen [+round] harmony is blocked by [+high] vowels and that blocking effect is illustrated in (116)-(117). When a [+high] vowel follows one of the rounding triggers [o] or [ɔ] the vowel to the right of that [+high] vowel is associated to [-round]. Unlike with ATR harmony in Akan, B. Orochen [+high] blockers do not have to share the [+round] feature of their predecessors; in other words [+round, +high] and [-round, +high] blockers are interchangeable. Lastly, the [+high] vowels do not exhibit blocking but transparency in words with [-round] vowels; [-round] vowels occur on either side of a [+high] vowel.

(116) [+high] vowels block [+round] harmony (Li 1996)

- | | |
|-------------------------------------|---|
| a. owon+dulə ‘pancake (destin)’ | d. ɔɾɔn+dɔla: ‘reindeer (destin)’ |
| b. bolboxi+wə ‘wild duck (def.acc)’ | e. tʃɔlik+pa ‘cloud-shaped design(def.acc)’ |
| c. moliktə ‘a kind of wild fruit’ | |

(117) a. [ɔɾɔndɔla:] ‘reindeer (destin)’ b. [bolboxi+wə] ‘wild duck (def.acc)’




The AR in (117a) illustrates that the domain of [+round] harmony is constrained by the leftmost [-high] feature. The first two vowels are [+round, -high, -long] and the third vowel is [+round, +high, -long]. While the [+high] blocker is also associated to the same [+round] feature as the preceding vowels it forms the right boundary of diffusion for that [+round] feature so the vowel to its right is associated to [-round]. There are also B. Orochen words with blocking vowels that are [-round, +high], such as in (116b) and (116d). The AR in (117b) illustrates one such word where the [-round, +high] blocker is associated to a different round feature than the vowels preceding it and the [-round] feature it associates to is also associated to the succeeding vowel.

The B. Orochen vowel inventory includes [+high] vowels that are also [+round], [u, ʊ]. Because these two vowels are neutral they can not have triggered rounding harmony in B. Orochen; on the surface a [-high] vowel to its right is associated to [-round]. Any [+round] vowel to the right of a [+high] vowel is also necessarily [+high, +round] [u] or [ʊ], as shown in (118h) and (119).

(118) Words beginning with [+round, +high] vowels (Li 1996)

- | | |
|---------------------|---|
| a. nuriktə 'hair' | f. urə+jə 'mountain (indef.acc)' |
| b. uʃi: 'rope' | g. luxi+xə:n+mə 'arrow (dim+def.acc)' |
| c. dʒu:xin 'otter' | h. imuksə+rək 'oil container (der.sfx)' |
| d. suxə 'axe' | |
| e. urə: 'earthworm' | |

(119) [imuksə+rək] 'oil container (der.sfx)'




The AR in (119) illustrates the iteration of [+round], but according to the generalization described thus far and because the two vowels associated to [+round] are also associated to [+high] I cannot claim that (119) exhibits B. Orochen round harmony. However, it is safe to say that B. Orochen utilizes diffuse [+round] harmony.

On the other hand, the behavior of initial [+round, +high] [u] and [ʊ] can be better understood when they are viewed alongside other [+high] vowels within words containing [-round] vowels. In short, [+high] vowels do not generally behave as blockers in words with [-round] vowels. Instead they are transparent on the surface because [-round] vowels occur on either side of a [+high] vowel, as shown in (120a-c) and (121) below.

(120) [+high] vowels are transparent to [-round] (Li 1996)

- | | |
|---------------------------------|-------------------------------------|
| a. bəjun+ksə 'elk hide' | d. bɪra+ja 'river (indef. acc.)' |
| b. əwulən 'the Big Dipper' | e. luxi+xə:n+mə 'arrow (def. acc.)' |
| c. əwəŋki+tʃien 'Evenki people' | f. bɪra+xə:n+ma 'river (def. acc.)' |
| g. bəjun 'elk' | i. əwi 'to play' |
| h. bəjə+ni 'person (poss.)' | |

(121) [əwulən] ‘the Big Dipper’




The examples in (120a-c) and the AR in (121) illustrate the iteration of [-round] across a [+high] vowel which has been shown to indicate transparency. The remaining examples in (120) show that [+high] vowels can also occur at the beginning of a word with [-round] vowels as in (118) and (120d-f) and at the end of such a word as in (120g-i). So the [+high] vowels exhibit blocking in words with [+round] vowels and transparency in words with [-round] vowels.

6.1.3 B. Orochen over multi-tiered ARs


In B. Orochen the round harmony pattern with [+high] blockers is describable using only the three FSCs over multi-tiered ARs in (122). The first FSC in (122a) describes a positional restriction on the cooccurrence of [+round, -high] vowels [o, ɔ, ɔ, ɔ:] with [+round, -high, +long] vowels [oɪ, ɔ:]. The FSC forbids a [o, ɔ, ɔ, ɔ:] from following an initial [o:] or [ɔ:] at any distance.

(122) FSCs for Baiyinna Orochen round harmony




The word in (123a) is grammatical because the initial [o:] is followed by only [-round, -high] vowels so it does not contain the FSC from (122a). However, if the second vowel was also associated to [+round] as in (123b) then the AR would contain the FSC and would thus be ungrammatical.

(123) a. [kɔ:məxə] ‘windpipe’ (Li 1996)




b. * [kɔ:məxə]




The word boundary symbol (#) is also needed in this FSC to specify that only initial [+round, -high, +long] vowels [o:, ɔ:] are forbidden from triggering [+round] harmony because in non-initial position they can still propagate it. In other words, on the surface an initial [o:] or [ɔ:] can only be followed by [-round, -high] vowels but in non-initial position a [o:] or [ɔ:] must have [+round, -high] vowels on either side of it. For example, in (124a) it is grammatical for a [+round, -high] [ɔ] to succeed a [+round, -high, +long] [ɔ:] that is word-medial.

(124) a. [ɔmɔ:ŋ+mɔ] (Li 1996)




b. *[ɔmɔ:ŋ+mɔ]




However, (124b) shows that if a [+round, -high, +long] [ɔ:] succeeds a word boundary symbol the AR with all [+round, -high] vowels is ungrammatical because it now contains the FSC from (122a).

The second FSC in (122b) forbids the iteration of both [+round] and [-round] features amongst [-high] vowels in B. Orochen. A [-round, -high] vowel [ə, ə:, a, a:] cannot be followed by a [+round, -high] vowel at any distance as long as both are associated to the same [-high] feature. This FSC describes both local and long distance assimilation because it includes connectedness on different tiers. It also applies even when other harmonizing vowels occur between the two forbidden vowels because the FSC does not include a successor relation between the vowels on the vocalic tier. In other words the successor relation that connects the features on the round tier along with the absence of such a connection on the vocalic tier describes the long distance round harmony pattern in a local way. For example, a word with an initial [ə, ə:, a, a:] cannot have any [+round, -high] vowels thereafter so (125a) does not contain (122b) and is thus grammatical.

(125) a. [əgdəgə] 'big' (Li 1996)




b. *[əgdogə]




But if the second vowel is associated to [+round] as in (125b) the AR now contains the FSC from (122b) and is ungrammatical. In short, the [-round] feature must be diffused: associated to all the [-high] vowels. This FSC also enforces [+round] diffusion in the same way by forbidding the alternation of round features, as shown in (126) below. In (126a) a word with full [+round] diffusion is grammatical as expected and does not contain (122b).

(126) a. [sokko: +mpo] ‘muddy(water) (contem)’ (Li 1996)



b. *[sokkə: +mpo]




However, in (126b) the medial [-round] [ə] causes the AR to contain the FSC from (122b) and thus be ungrammatical. So whether it is in initial position or not (122b) forbids a [-round] vowel from being followed by a [+round] vowel when both are associated to the same [-high] feature. Even with the arrow on the round tier (122b) enforces the diffusion of both [+round] and [-round] amongst [-high] vowels. In addition, the


diffusion on the high tier allows words to contain vowels associated to different round features when a [+high] vowel intervenes. Examples and discussion of the [+high] blocking vowels will come later.

The FSC in (122c) forbids a word which contains a vowel associated to [+round, -high, -long] [o, ɔ] from being followed by a [-round, -high] vowel [ə, ə:, a, a:] when both vowels are associated to the same [-high] feature. This FSC is almost the reverse of (122b) except that it includes a [-long] feature associated to the leftmost vowel. This extra [-long] feature is necessary to allow grammatical words which begin with a [+round, -high, +long] [o:, ɔ:] followed by [-round, -high] vowels like in (123a). The example in (127a) shows that an AR with a fully diffuse [+round] feature is grammatical when the leftmost vowel is associated to all three of [+round, -high, -long].

(127) a. [sokko:mpɔ] ‘muddy(water) (contem)’ (Li 1996)



b. *[sokko:mpə]




The AR in (127b) illustrates that an AR which contains a noninitial [+long] vowel can violate the FSC in (122c) if it does not include full diffusion of [+round]. The final [-round, -high] vowel causes this AR to be ungrammatical because it contains the forbidden substructure from (122c). In particular, this FSC includes the successor relation on only the round tier so the AR is ungrammatical because the final vowel is associated to [-round, -high] even though a [+long] feature intervenes.


The connectivity on the high tier of (122c) also describes the local aspect of blocking in the B. Orochen round harmony pattern. The FSC specifically forbids two vowels from being associated to different round features when they are both associated to a single [-high] feature. This configuration allows for words with [+high] blockers because while [-high] vowels on either side of a blocker are associated to different round features they are also associated to different iterations of a [-high] feature, as shown in (117a) and repeated below in (128a).

(128) Blocking in Baiyinna Orochen

- a. [ɔrɔndɔla:] ‘reindeer (destin)’ (Li 1996)



- b. *[ɔrɔndɪɛla:]



The AR in (128b) is ungrammatical because the diphthong is associated to [-round] and to the same single [-high] feature that is associated to the preceding [+round, -long] vowel. So while the [+high] feature in (128a) allows the AR to contain both [+round] and [-round] vowels, replacing the [ʊ] with a [-high] diphthong illustrates how a fully diffused [-high] feature does not allow a [-round] feature to succeed a [+round] feature because the AR in (128b) contains the forbidden substructure of (122c). So even though high and round features are represented on separate unrelated tiers they can interact as a result of their associations to vowels on a single vocalic tier.


On the surface, B. Orochen round harmony is represented by the diffusion of [+round] amongst all [-high] vowels. When [o] or [ɔ] occur after a [-round, +high] vowel all of the [-high] vowels to the right of that

[-round, +high] vowel are associated to [+round]. The FSC in (122b) enforces that the vowel which precedes the leftmost [+round, -high] must be associated to [+high]. Examples of the diffusion of B. Orochen round harmony after an initial [-round, +high] vowel are found in the non-native borrowings shown in (129) below.

(129) round harmony with non-initial [+round, -high] vowels (Li 1996)

- a. kinɔ+wɔ ‘film (def.acc) (Russian)’
- b. dʒijgon+lɔ:tʃo ‘to attack (der.sfx+pst) (Chinese)’
- c. gʊaŋbɔ+lɔ:tʃɔ ‘to broadcast (der.sfx+pst) (Chinese)’

(130) [kinɔwɔ] ‘film (def.acc) (Russian)’



The AR of (129a) is shown in (130) and illustrates the multiple association of [+round] to the rightmost two [ɔ] vowels, excluding the leftmost vowel; longer words with additional suffixes demonstrate the same pattern. Crucially, the diffusion of both [+round] and [-high] features does not contain the FSC from (122c) and the [+high] feature associated to the initial [-round] vowel prevents the AR from containing the FSC from (122b).¹⁰

In this section I have shown that B. Orochen is describable with a CNL of only three FSCs over multi-tiered ARs. In the next section I will present my analysis of B. Orochen over strings of vowels.

6.1.4 B. Orochen over strings

Baiyinna Orochen vowel harmony can also be analyzed as a set of vowel strings consisting of the vowels in the inventory described above. The B. Orochen stringset \mathbb{L}_{BO} is summarized in Table (22) below. Since the data makes it clear that round harmony applies both within roots and across morpheme boundaries in B. Orochen I do not include morpheme boundaries in the vowel strings in this section.

¹⁰The word in (129c) includes root disharmony that is not seen in any other examples and is most likely present here because this word is borrowed from Chinese. Despite the preceding [-round, -high] [a] the [+round] feature is still diffused across morpheme boundaries and would be associated to all three succeeding vowels which suggests that [+round] harmony has applied from the root-final [+round, -high, -long] trigger.

Table 22: Baiyinna Orochen Vowel Strings

	strings $\in \mathbb{L}_{BO}$	strings $\notin \mathbb{L}_{BO}$
	ɪ, uə, ə:a	io, u:ɔ, ɔa
Harmony:	oo, ɔɔ:, ɔ:a	o:ɔ, aɔ:, ɔə
	ooo, oo:ɔ, ɔɔ:	iuo, oəo:, əɔa
Blocking:	ouə, oia:, ɔa	oio:, ɔɔ:ɔ, ɔɔ:

This stringset includes strings with only [-round] vowels, only [+round] vowels, or both when a [+high] vowel intervenes. It also includes strings with multiple [+high] blockers.

The B. Orochen stringset \mathbb{L}_{BO} is SL_3 because it can be described using a CNL with restrictions of length two and three and these restrictions are interpreted with the successor relation. I have broken up the CNL in (131) into five chunks in order to clearly explain how each set of NLS corresponds to a different crucial aspect of the B. Orochen round harmony pattern I have described so far. The featural NLS serve only as a summary to help clarify the pattern described by the larger set of vocalic NLS. In addition, I included the full CNL below with both [+ATR] and [-ATR] vowels in order to illustrate the immensity and thoroughness needed to describe the B. Orochen round harmony pattern; but this language also utilizes ATR harmony. ATR harmony would eliminate any words with combinations of [+ATR] and [-ATR] vowels from \mathbb{L}_{BO} .

(131) SL_3 CNL for \mathbb{L}_{BO} , interpreted with successor ($\langle \rangle$)

\mathbb{L}_{BO} :

- a. $\neg\#[+round, -high, +long][+round, -high] = \neg\#o:ɔ \wedge \neg\#o:ɔ: \wedge \neg\#o:ɔ \wedge \neg\#o:ɔ: \wedge \neg\#ɔ:ɔ \wedge \neg\#ɔ:ɔ:$
 $\wedge \neg\#ɔ:ɔ \wedge \neg\#ɔ:ɔ: \wedge$
- b. $\neg[+round, -high, -long][-round, -high] \wedge \neg[-round, -high][+round, -high] = \neg oə \wedge$
 $\neg ɔə \wedge \neg oə: \wedge \neg ɔə: \wedge \neg oə \wedge \neg ɔə \wedge \neg oə: \wedge \neg ɔə: \wedge \neg əo \wedge \neg əo: \wedge \neg əo \wedge \neg əo: \wedge$
 $\neg ə:ɔ \wedge \neg ə:ɔ: \wedge \neg ə:ɔ \wedge \neg ə:ɔ: \wedge \neg aɔ \wedge \neg aɔ: \wedge \neg aɔ \wedge \neg aɔ: \wedge \neg a:ɔ \wedge \neg a:ɔ: \wedge \neg a:ɔ \wedge$
 $\neg a:ɔ: \wedge$
- c. $\neg[+round, -high, -long][+round, -high, +long][-round, -high] = \neg oo:ə \wedge \neg oo:ə: \wedge \neg oo:ə \wedge \neg oo:ə: \wedge$
 $\wedge \neg oo:ə: \wedge \neg oo:ə: \wedge \neg oo:ə: \wedge \neg oo:ə: \wedge \neg oo:a \wedge \neg oo:a: \wedge \neg oo:a \wedge \neg oo:a: \wedge$
 $\neg oo:a: \wedge \neg oo:a: \wedge \neg oo:a: \wedge \neg oo:a: \wedge$
- d. $\neg[+round, -high, +long][+round, -high, +long][-round, -high] = \neg o:ɔ:ə \wedge \neg o:ɔ:a \wedge$
 $\neg ɔ:ɔ:ə \wedge \neg ɔ:ɔ:a \wedge \neg o:ɔ:ə \wedge \neg o:ɔ:a \wedge \neg ɔ:ɔ:ə \wedge \neg ɔ:ɔ:a$

The first set of NLs in (131a) are restrictions of length $k = 3$ which begin with a word boundary symbol (#). These NLs describe which vowels cannot occur directly succeeding a word boundary; in other words which vowels cannot occur at the beginning of a B. Orochen word. In the first line of vocalic NLs we see that vowel strings in \mathbb{L}_{BO} cannot begin with a [+round, -high, +long] vowel [o:] or [ɔ:] which is succeeded by a [+round, -high] vowel [o, ɔ, ɔ, ɔ:]. So in (132a) the word [kɔ:ŋakta] meaning ‘handbell’ is grammatical because in its vowel string the initial [+round, -high, +long] [ɔ:] is succeeded by a [-round, -high] [a] so it does not contain any of the NLs in (131a).

(132) [kɔ:ŋakta] ‘handbell’ (Li 1996)

a.
→ ɔ: → a → a ¬ # → ɔ: → ɔ

b. * [kɔ:ŋokta]

→ ɔ: → ɔ → a ¬ # → ɔ: → ɔ

The hypothetical word in (132b) is ungrammatical because the second vowel is [+round, -high] so the vowel string contains the NL from (131a). The first eight NLs thus describe how [o:] and [ɔ:] do not participate in [+round] harmony from a word-initial position.

The next set of NLs in (131b) are only length $k = 2$ and describe the basic fact of rounding harmony in B. Orochen: [+round, -high] [o, ɔ, ɔ, ɔ:] and [-round, -high] vowels [ə, ə:, a, a:] cannot succeed one another. More specifically, the first eight NLs state that a [+round, -high, -long] vowel [o, ɔ] cannot be succeeded by a [-round, -high] vowel [ə, ə:, a, a:]. For example, in (133a) the vowel string is grammatical because all the vowels are [+round, -high] illustrating full [+round] harmony; it does not contain any of the NLs from (131b).

(133) [ɔlɔ+jɔ] ‘fish (indef. acc)’ (Li 1996)

a.
ɔ → ɔ → ɔ ¬ ɔ → a

b. * [ɔla+jɔ]

ɔ → a → ɔ ¬ ɔ → a

b. * [olo: + nə + tfo]

o → o: → ə → o ¬ o → o: → ə

However, (135b) illustrates how changing the third vowel to be [-round, -high] makes the vowel string and hypothetical word ungrammatical because the non-initial [o:] does not block [+round] harmony. Since there are no word boundaries in these NLs a string which begins with more than one [+round, -high, -long] vowel [o, ɔ] will still be ungrammatical if it contains the other two vowels in succession thereafter. The previous two sets of NLs show that vowel strings with initial [+round, -high, +long] [o:, ɔ:] do not exhibit [+round] diffusion and this set of NLs show that those same vowels do not block [+round] harmony in non-initial positions.

The fourth set of NLs in (131d) are also of length $k = 3$ and they eliminate a crucial set of vowel strings which would be predicted to be ungrammatical in B. Orochen. These NLs state that no B. Orochen vowel string can contain a [+round, -high, +long] [o:, ɔ:] succeeded by another of the same which is succeeded by a [-round, -high] [ə, ə:, a, a:]. The hypothetical word in (136a) is a grammatical extension of the word from (135a). The hypothetical vowel string below is grammatical because it exhibits full [+round] harmony and thus does not contain any of the NLs from (131d).

(136) [olo:ɔ:o: + no + tfo] (Li 1996)

a.

o → o: → o: → o: → o → o ¬ o: → o: → ə

b. * [olo:ɔ:o: + nə + tfo]

o → o: → o: → o: → ə → o ¬ o: → o: → ə

The hypothetical word in (136b) on the other hand is ungrammatical no matter how many [o:] vowels there are because there is no rounding blocker so having the [-round, -high] [ə] amongst a long string of [+round, -high] vowels means that string contains one of the NLs from (131d). Just as the previous set of NLs ruled out strings with [+round] vowels succeeded by [-round] vowels regardless of how many might occur at the beginning so too does this set rule out strings with [+round] vowels succeeded by [-round] vowels regardless of how many [+round] vowels might occur in the middle. In other words, the 3-factors which are forbidden by this final set of NLs can occur anywhere within a longer vowel string and that vowel string will be ungrammatical. In conjunction with the previous set of NLs these eight NLs effectively enforce [+round] harmony at any

distance and rule out the possibility of having more than one [+round, -high, +long] [o:, ɔ:] at the beginning of a vowel string.

I have shown that B. Orochen contains both a long distance [+round] harmony and a local blocking pattern. B. Orochen [+round] harmony with blocking also relies on vowel length which does not usually interact with segmental patterns like vowel harmony. This pattern is describable with only three FSCs over multi-tiered ARs which is significantly fewer than the 56 needed to describe it over vowel strings.

6.2 Kinande

Kinande is a Bantu language spoken in the Democratic Republic of Congo and Uganda. Kinande utilizes 10 vowels distinguished by height and ATR features. The [+high] vowels are [i, ɪ, u, ʊ], [+low] vowels are [ə, a], and [-high, -low] vowels are [e, ɛ, o, ɔ]. The [+ATR] vowels are [i, u, e, o, ə] and the [-ATR] vowels are [ɪ, ʊ, ɛ, ɔ, a]. The crucial feature distinctions for ATR harmony among Kinande vowels are as follows: [+high, +ATR] vowels are [i, u], [+high, -ATR] vowels are [ɪ, ʊ]; [-high, -low, +ATR] vowels are [e, o], [-high, -low, -ATR] vowels are [ɛ, ɔ]; the [+low, +ATR] vowel is [ə] and the [+low, -ATR] vowel is [a]. The Kinande vowels and their respective features relevant to the harmony pattern discussed in this section are shown in Table 23 (Gick et al. 2006; Mutaka 1995).

Table 23: Kinande Vowels

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

The early literature on Kinande included some conflict between analyses of Kinande [+low] vowels, which was then clarified with acoustic and ultrasound evidence. Mutaka (1995) relies on the author's native speaker judgements to claim that Kinande uses a single phonologically low vowel. Rather than acknowledging a difference between the low vowels that occur in different environments he suggests that a single [+low] vowel is transparent to ATR harmony. Hyman (2002) did not take a clear stance on the status of the [+low] vowel [a] when he claimed that Kinande utilizes ATR harmony in both directions so [a] can be either transparent or undergo harmony in one direction and it is opaque in the other direction. Archangeli and Pulleyblank (2002) provided an alternative analysis of ATR harmony in Kinande based on the following assumptions:

- (a) the Kinande vowel inventory looks like Table 23, with two low vowels distinguished by a binary ATR feature and

(b) low vowels participate in vowel harmony rather than being transparent to it.

Gick et al. (2006) then conducted a series of acoustic and ultrasound studies, which confirmed Archangeli and Pulleyblank (2002)'s assumption (a) with phonetic evidence that Kinande uses two phonologically low vowels and those low vowels can be productively distinguished using a binary categorical ATR feature. Gick et al. (2006) further agreed with Archangeli and Pulleyblank (2002) by stating that they expect assumption (b) to be true. My analysis follows both assumptions (a) and (b) and as throughout this dissertation I rely on the binary categorical ATR feature as Gick et al. (2006) suggest.

The basic generalization for Kinande ATR harmony can be described as follows: on the surface all [-high] vowels have the same ATR feature as the closest [+high] vowel to their right. In the absence of any [+high] vowels the surface form of a word contains only [-ATR] vowels. Some examples of full harmony with [+ATR] and [-ATR] vowels are shown in (139) below. This section uses data cited with the following abbreviations: (M) for Mutaka (1995), (H) for Hyman (2002), and (AP) for Archangeli and Pulleyblank (2002).

(137) Full ATR Harmony

<u>+ATR</u>	<u>-ATR</u>
a. $\varepsilon + rili:ba$ 'to cover' (AP)	f. $ka:ra$ 'force for/at' (H)
b. $tute\beta inira$ 'we are not dancing for' (M)	g. $\varepsilon + rɪlangɪr$ 'to see' (M)
c. $m\alpha + tw\acute{a}kiheki:r\acute{e}$ 'we carried it' (H)	h. $\varepsilon + rɪlɪm:r\acute{a}$ 'work (for)' (M)
d. $\varepsilon + kik\acute{e}:li$ 'woman' (M)	i. $ama:tu\acute{g}u$ 'yams' (H)
e. $m\alpha + muk\acute{a}mi$ 'brewer' (AP)	j. $m\alpha + twasw\acute{e}:r\acute{e}$ 'we ground' (M)

The data in (137) support this description of Kinande ATR harmony. In the right column (137f-j) contain only [-ATR] vowels, and in the left column there is more variety. The example in (137j) illustrates the behavior described above: no [+high] vowels and all vowels are [-ATR]. In (137b-c) the rightmost [+high] vowel is [+ATR] and each word ends with a [-ATR] vowel. This pattern can be explained by claiming that Kinande ATR harmony affects vowels to the left of *and not to the right of* a [+high] vowel. The [+high] vowels thus appear to block the diffusion of [+ATR].

The Kinande generalization might appear to predict that words with a final [+low] vowel to the right of the rightmost [+high] vowel could end in either [a] or [ə] because ATR harmony does not affect [-high] vowels to the right of a [+high] vowel. However, blocking requires that different values of the harmonic feature occur on either side of a blocker. Since [-ATR] vowels surface on the right and [+ATR] vowels surface on the left of [+high] vowels it looks like [+high] vowels block ATR harmony. But the same is not true in the reverse order;

in (137f-g) the vowels on both sides of a [+high] vowel are [-ATR] as is the [+high] vowel so [-ATR] is fully diffused in these examples despite the fact that they include a [+high] vowel.

In addition, the very first vowel in (137a, c-e, g-h, j) is [-ATR]. The fact that these begin with a [-ATR] vowel appears to flaunt the generalization I just described. However, Mutaka (1995) explains that the initial vowel in Kinande belongs to a particular morpheme—[ɛ] or [ɔ̃]—which is added after the phonological ATR harmony pattern has applied. In short, the initial vowel—[ɛ] or [ɔ̃]—is always treated as outside the domain of ATR harmony. This stipulation holds true for all Kinande words with either of these prefixes and so while the data is written as it appears in the sources, ARs will exclude these two initial vowels since they do not participate in ATR harmony. I also use the morpheme boundary symbol above to separate these prefixes from the rest of the harmony domain which can include multiple morphemes.

Lastly, Archangeli and Pulleyblank (2002) claim that Kinande ATR harmony is constrained by morphological domains, but I propose a purely phonological one. On the surface, ATR harmony applies to all [-high] vowels to the left of a [+high] vowel and each [+high] vowel denotes a new harmony domain.

(138) Multiple harmony domains within a word

- a. ɛ+rilibanira ‘disappear (for someone)’ (M)
- b. ɛ+rihimatira ‘to press for/at’ (H)
- c. ɛ+rigumatra ‘to stuff in mouth for/at’ (H)

Thus longer words with multiple [+high] vowels like in (138) can include multiple ATR harmony domains, each with a different ATR feature value. In short, a Kinande word can include multiple ATR features as long as each [-high] vowel has the same ATR feature value as the closest [+high] vowel to its right.

The remainder of this section analyzes the computation of the Kinande generalization described here. Section 6.2.1 gives my analysis of Kinande ATR harmony over multi-tiered ARs. Section 6.2.2 gives my analysis of Kinande ATR harmony over strings of vowels. And section 6.2.3 concludes.


6.2.1 Kinande over multi-tiered ARs

In this section I analyze this pattern over multi-tiered ARs to show that it is ASL^{VH}. Kinande ATR harmony can be broken down into three separate concepts: diffusion of [+ATR] when associated to a [+high] vowel, diffusion of any ATR feature amongst [-high] vowels, and the full diffusion of [-ATR] when no [+high] vowels are present. Over multi-tiered ARs the CNL contains three types of FSCs which correspond to each of these concepts. In this section while the Kinande words are written out in full, ARs exclude the initial [ɛ] or [ɔ̃] vowel because it is excluded from the harmony domain. These prefixes are separated from the harmony domain by a morpheme boundary symbol.


The ARs for the two words shown in (139) exemplify the Kinande ATR harmony generalization. The example in (137a) contains all [+ATR] vowels within the domain of harmony and (137h) is identical except that it contains only [-ATR] vowels; each word has a different and unrelated gloss.

(139) Kinande Basic ATR Harmony ARs

a. [ε†rili:ba] ‘to cover’ (AP)



b. [ε†rɪlɪmɪ:ra] ‘work (for)’ (M)



The surface AR in (139a) illustrates the [+ATR] harmony generalization described above. While each [+high] vowel could introduce a different ATR harmony domain, they are all associated to [+ATR] and the final vowel is associated to [-ATR] because it occurs to the right of the rightmost [+high] vowel. The surface AR in (139b) exhibits full [-ATR] diffusion.

Examples like those in (138) and (140) illustrate multiple ATR harmony domains. These examples contain [+high] vowels associated to different ATR features within a single word and so there are [-high] vowels associated to ATR features with different values. These are grammatical because each [+high] vowel denotes a new ATR harmony domain.


(140) Multiple harmony domains within a word

a. ε†rilibanɪra ‘disappear (for someone)’ (M)

b. ε†rihimatra ‘to press for/at’ (H)

c. ε†rigumatra ‘to stuff in mouth for/at’ (H)

(141) [εrilibanɪra] ‘disappear (for someone)’




For example, the AR in (141) contains a single [-ATR] feature associated to the rightmost vowel, the rightmost [+high] vowel, and the antepenultimate vowel. However, further leftward the first two [+high] vowels are

associated to a single [+ATR] feature. This word thus appears to contain two harmony domains: the rightmost [+high] vowel falls within a series of vowels associated to [-ATR], but the preceding two [+high] vowels are associated to [+ATR]. Without any other evidence I would predict that a final vowel to the right of a [+high] vowel could be associated to either [+ATR] or [-ATR]; but Kinande also includes words with no [+high] vowels at all, such as *mɔtʰwaswɛrɛ* ‘we ground’, and words with a [-ATR, +high] vowel such as *karra* ‘force for/at’; these are discussed in more detail below. In such words all the vowels are associated to [-ATR].

All the facts of Kinande ATR harmony that have been laid out so far can also be described using the five FSCs over multi-tiered ARs in (142). FSCs over multi-tiered ARs also illustrate how the Kinande ATR harmony pattern can be broken down into three different subpatterns: diffusion of [+ATR] to the left of a [+high] vowel, diffusion of ATR features amongst [-high] vowels, and the diffusion of [-ATR] on the right of a [+high] vowel.



(142) Kinande FSCs




The first two FSCs in (142a)-(142b) describe the diffusion of an ATR feature which is associated to a vowel that is also associated to [+high]. In short, a [+high] vowel cannot succeed a [-high] vowel which is associated to an ATR feature with a different value. For example the AR of the Kinande word in (143a) below

is grammatical because it does not contain the FSC from (142a) or (142b); while the AR does contain two successive ATR features with different values, the vowel associated to [-ATR] is also associated to [-high] and the [+high] feature is associated to a preceding vowel which is associated to a diffuse [+ATR] feature. In other words the AR in (143a) is grammatical because the [+ATR] feature is diffused and the rightmost [+ATR] vowel is also associated to [+high].

(143) a. [motwəsoli:rɛ] ‘we paid taxes’ (M)



b. *




Unlike in other constraints throughout this dissertation (142a-b) include the successor relation on the vowel tier. The ungrammatical AR in (143b) illustrates exactly what this extra successor relation accomplishes. Including the successor relation on the vowel tier makes (143b) ungrammatical because the second and third vowels are associated to different ATR features; but if the initial vowel was associated to [-ATR] rather than the second vowel the AR would not violate this FSC. The same holds true for words with [-ATR, +high] vowels which are restricted by the FSC in (142b). For other patterns I was able to describe long distance harmony by only including successor relations on the feature tiers within the FSCs. However, Kinande words can include multiple ATR harmony domains when there is more than one [+high] vowel so these first two

FSCs only describe the diffusion of ATR within a single domain. In order to prohibit disharmony over longer distances within a single domain I conjoin (142a-b) with another set of FSCs discussed in more detail below.

The second two FSCs in (142c-d) describe the diffusion of an ATR feature amongst [-high] vowels. No vowel can succeed a vowel associated to an ATR feature with a different value if both are associated to [-high]. As in (143a) the AR in (144a) below is grammatical because the [-ATR] feature is diffused across successive vowels which are all associated to [-high] and so it does not contain the FSC from either (142c) or (142d). The AR in (144a) illustrates how constraints which refer only to the successor relation (even between vowels) can still have long distance effects; a single ATR harmony domain will contain only [-high] vowels that are associated to the same ATR feature as the [+high] vowel to their right because no two successive [-high] vowels can be associated to different ATR features.

(144) a. [mɔtwaswɛ:rɛ] ‘we ground’ (M)



b. *




The AR in (144b) illustrates how the two FSCs in (142c-d) describe this fact; there is no [+high] feature and it is ungrammatical because two successive vowels are both associated to [-high] but associated to different ATR features. Thus no matter how long a vowel string with no [+high] vowels is the [-high] vowels must all be associated to a single [-ATR] feature.

This brings me to the final FSC in (142e) which forbids a vowel from being associated to both a final [-high] and a final [+ATR] feature. This FSC enforces that without a [+high] vowel to its right [-ATR] will be diffused across all [-high] vowels up to a [+high] vowel to the left. When conjoined with (142c) and (142d), (142e) essentially forbids a vowel string with no [+high] feature from having a [+ATR] feature. The final boundary is represented by the symbol on the right of the FSC (#).


(145) a. [mɔtwaswɛ:rɛ] ‘we ground’ (M)



b. [mɔtwəsɔli:rɛ] ‘we paid taxes’ (M)



c. *



The grammatical AR in (145a) contains only a [-ATR] feature and so it cannot contain the FSC from (142e). The grammatical AR in (145b) includes both a [+ATR] and a [-high] feature, but they are not associated to the same vowel and only the [-high] feature is succeeded by a final boundary symbol so it cannot contain the FSC from (142e). As I mentioned earlier, without evidence from examples like (145a) I would predict that final [-high] vowels to the right of a [+ATR, +high] vowel could be associated to either [+ATR] or [-ATR]. The FSC in (142e) contradicts that prediction by allowing (145a) and forbidding (145c).

This final FSC forbids an AR in which a vowel associated to [-high] is also associated to [+ATR] when both features are succeeded by a final boundary. The use of a final boundary symbol (#) differs from all the other FSCs proposed in this dissertation so far. My previous work in Blum (2019) showed that adding this string-edge boundary symbol to the alphabet for FSCs over multi-tiered ARs can allow such a CNL to describe patterns like sour grapes and I argued that a restrictive theory of phonology should not be expressive enough to

describe unattested patterns. However, recent work has shown that sour grapes may not be as unattested as previously thought (McCollum et al. 2020) and so including string-edge boundaries could still be an acceptable addition to the set of symbols allowed on a tier. In short, the expressive power of the string-edge boundary symbol deserves further investigation, but such an investigation is beyond the scope of this dissertation.

In this section I have shown that Kinande is describable with a CNL of only five FSCs over multi-tiered ARs. In the following section I present my analysis of Kinande over strings of vowels.

6.2.2 Kinande over strings

The Kinande ATR harmony pattern can also be described as follows: [-high] vowels to the left of a [+high] vowel have the same ATR feature as that [+high] vowel and without any [+high] all vowels are [-ATR]. If this pattern makes up the stringset \mathbb{L}_K then a sample of the sets of strings with lengths $k=2$ and $k=3$ which are members of the stringset \mathbb{L}_K and those which are not members of \mathbb{L}_K are listed below.

Table 24: Kinande Vowel Strings

	strings $\in \mathbb{L}_K$	strings $\notin \mathbb{L}_K$
No final [+high]:	$\epsilon\circ, i\epsilon, \circ a\epsilon, u\epsilon a \dots$ $ueiia, \circ iei\epsilon, \dots$	$\circ e, i\circ, ai\epsilon, \circ i\circ, \dots$ $uaiia, ai\epsilon i\epsilon, \dots$
With final [+high]:	$\circ i, \circ u, \epsilon i, a\circ, \dots$ $\epsilon \circ i, i\circ u, \epsilon u i, \epsilon a\circ, \dots$ $iia i, aa\circ\circ, \dots$	$\epsilon i, \circ i, \circ u, \epsilon\circ, \dots$ $\epsilon \circ i, u\circ i, a\epsilon u, \circ a\circ, \dots$ $iia i, a\circ\circ u$

The ATR harmony pattern over strings described by the generalization above is Strictly Local (SL_2) because it is closed under suffix substitution and it is describable with the CNL of FSCs over vowel strings of length $k=2$ in (146) interpreted with the successor ordering relation. As before the featural CNL is provided solely to help clarify the pattern described by the vowel string CNL.

(146) SL CNL for \mathbb{L}_K

$\mathbb{L}_K =$

a. $\neg[-ATR, -high][+ATR, +high] \wedge \neg[+ATR, -high][-ATR, +high] = \neg \epsilon i \wedge \neg \epsilon u \wedge \neg \circ i \wedge$

$\neg \circ u \wedge \neg ai \wedge \neg au \wedge \neg \epsilon i \wedge \neg \epsilon u \wedge \neg \circ i \wedge \neg \circ u \wedge \neg \circ i \wedge \neg \circ u \wedge$

b. $\neg[+ATR, -high][-ATR, -high] \wedge \neg[-ATR, -high][+ATR, -high] = \neg \epsilon \epsilon \wedge \neg \epsilon \circ \wedge \neg \epsilon a \wedge$

$\neg \circ \epsilon \wedge \neg \circ \circ \wedge \neg \circ a \wedge \neg \circ \epsilon \wedge \neg \circ \circ \wedge \neg \circ a \wedge \neg \epsilon \epsilon \wedge \neg \epsilon \circ \wedge \neg \epsilon \circ \wedge \neg \circ \epsilon \wedge \neg \circ \circ \wedge \neg \circ \circ \wedge \neg \circ a \wedge \neg \circ a \wedge \neg$
 $a\circ \wedge$

c. $\neg[+ATR, -high]\# = \neg \epsilon\# \wedge \neg \circ\# \wedge \neg \circ\#$

The CNL above illustrates that the Kinande ATR harmony pattern can be broken down into three concepts: [-high] vowels to the left of [+ATR, +high] vowels are also [+ATR], a span of [-high] vowels has the same ATR feature value, and without a [+high] vowel all vowels are [-ATR]. The NLs in (146a) describe the ATR behavior of [+high] vowels by forbidding [-high] vowels from being succeeded by a [+high] vowel with a different ATR feature. For example the grammatical string in (147a) does not contain any of the forbidden substructures in (146a). Two relevant NLs are shown on the right in bold and red.

(147) a. [mɔ+twəkiheki:ɛ] ‘we carried it’ (H)

ə → i → e → i → ε **¬ ei ∧ ¬ ai**

b. * [mɔ+twəkiheki:ɛ]

ə → i → **e → i** → ε **¬ ei**

c. * [mɔ+twakiheki:ɛ]

a → i → e → i → ε **¬ ai**

In (147b-c) a single [-high] vowel has been changed to [-ATR] but is still succeeded by a [+ATR, +high] vowel. So each example contains one of the NLs from (146a) shown in bold and red. The examples in (147) illustrate the fact that on the surface only a [-high] vowel which is also [+ATR] can be succeeded by a [+ATR, +high] vowel because the [+high] vowel denotes a domain of [+ATR] harmony.

The NLs in (146b) describe the harmonic behavior of [-high] vowels by forbidding them from being succeeded by a [-high] vowel with a different ATR feature. For example the grammatical word in (148a) contains all [-ATR, -high] vowels and does not contain any of the NLs from (146b). Two relevant NLs are shown on the right in bold and red.

(148) a. [mɔ+twəswɛ:ɛ] ‘we ground’ (M)

a → ε → ε **¬ əɛ ∧ ¬ ɛɛ**

b. * [mɔ+twəswɛ:ɛ]

ə → ɛ → ε **¬ əɛ**

c. * [mɔ+twəswɛ:ɛ]

a → **ɛ → e** **¬ ɛɛ**

In (148b) the first vowel is changed to its [+ATR] counterpart and in (148c) the final vowel is changed to its [+ATR] counterpart. Each of these examples is ungrammatical because they contain [+ATR, -high] vowels when there is no [+ATR, +high] vowel. Thus they each also contain one of the NLS from (146b) shown in bold and red.

And lastly, in conjunction with the FSCs in (146b) the final three NLS in (146c) use the string boundary symbol (#) to describe the behavior of [-ATR] vowels by forbidding [+ATR, -high] vowels from being succeeded by a word boundary. Since there is no [+high] vowel to the right of a final [-high] vowel the final vowel must be [-ATR]. In (149a) below the grammatical string contains a final [-ATR, -high] [a] and not any of the three NLS from (146c).

(149) a. [tuteβinira] ‘we are not dancing for’ (M)

u → e → i → i → a → # **¬ ə#**

b. *[tuteβinirə]

u → e → i → i → **ə** → # **¬ ə#**

The example in (149b) on the other hand has changed the final vowel to be the [+ATR] counterpart of [a]. Without a [+high] trigger to its right this final vowel is ungrammatical and thus the vowel string contains one of the NLS from (146c) in bold and red.

The Kinande ATR harmony pattern also illustrates another common aspect of vowel harmony: domains. In Kinande the [+high] vowels demarcate the right boundary of ATR harmony regardless of the ATR feature value associated with vowels to its right. Thus a Kinande vowel string can contain [-high] vowels with different ATR feature values, but they must be separated by a [+high] vowel. For example, a grammatical word like eriliβan:r:a ‘to disappear for someone’ contains multiple different ATR harmony domains indicated by the ATR values of the [-high] vowels.

(150) [eri+liβan+r:a] ‘to disappear for someone’ (M)

(e → i) → (i) → (a → r:) → a

The vowel string in (150) explicitly notates the separate domains present in this word, indicated with parentheses. First the rightmost vowel is [-high] and so without a [+high] vowel to its right the final vowel is [-ATR]. The penultimate vowel is [-ATR, +high] which denotes a [-ATR] harmony domain and so the adjacent preceding vowel is also necessarily [-ATR]. The post-penultimate vowel [i] is [+ATR, +high] and so it would denote a new [+ATR] harmony domain but there is no [-high] vowel directly preceding it to show this. However, the penultimate vowel is also [+ATR, +high] which denotes another [+ATR] harmony domain and requires

the initial vowel to be [+ATR]. You may notice that these vowel harmony domains do not line up with the morpheme boundaries and this is because the phonological vowel harmony pattern applies across morphemes in Kinande.

The Kinande pattern is also SL_2 because it does not involve transparency as Mutaka (1995) originally claimed. Chapters 4 and 5 showed that transparency is not SL_k and sometimes can even be more complex than other vowel harmony patterns. If Kinande utilized a single transparent [+low] vowel as Mutaka (1995) suggested then the pattern would be at least SP_k like Finnish but not SL_2 as shown above. A full analysis of the computation of Mutaka (1995)'s theory is beyond the scope of this dissertation. However, my analysis is supported by phonetic evidence for the vowel inventory of Kinande and my methodology reveals the computational simplicity of a pattern that previously required extra theoretical machinery to accurately describe. I have shown that the Kinande ATR harmony pattern is SL_2 over strings of vowels and that it can be described more succinctly using FSCs over multi-tiered ARs.

In this section I have shown that the Kinande ATR harmony pattern is SL_2 over strings and it is describable with a CNL of only five FSCs over multi-tiered ARs so it fits within the ASL^{VH} class of ARsets. The CNL over strings of vowels contains 33 FSCs while the CNL over multi-tiered ARs contains only five FSCs. Thus Kinande provides another example of how multi-tiered ARs allow us to describe SL patterns with fewer restrictions.

6.3 Conclusion

This chapter presents evidence for my second argument in favor of using of multi-tiered ARs: they provide more concise pattern descriptions with fewer FSCs. I discussed two vowel harmony patterns: Baiyinna Oroquen round harmony and Kinande ATR harmony which both fit into the SL_k stringset class. Both Baiyinna Oroquen and the Kinande also fit within the ASL^{VH} class of multi-tiered ARsets because they can be described with a CNL of FSCs over multi-tiered ARs. I also showed that over vowel strings the Baiyinna Orochen round harmony CNL requires 56 FSCs and the Kinande ATR harmony CNL requires 33 FSCs. On the other hand, over multi-tiered ARs Baiyinna Orochen round harmony is describable with a CNL of only three FSCs and Kinande is describable with a CNL of only five FSCs. So while multi-tiered ARs do not reduce the complexity of SL_k patterns they do provide a useful reduction in the number of FSCs needed to describe such patterns.

7 Conclusion

In this dissertation I have shown that multi-tiered autsegmental representations (ARs) both reduce the complexity of non-local vowel harmony patterns and provide concise descriptions of local vowel harmony patterns. I argue that for these two reasons ARs are a preferable abstract representation of vowel harmony as compared with strings of vowels. I thus propose the Autosegmental Strictly Local (ASL^{VH}) class of ARsets which is defined as containing vowel harmony patterns that are describable with a conjunction of negative literals (CNL) consisting of connected forbidden substructure constraints (FSCs) and are thus local over ARs. Being able to describe patterns in a local way is useful because it reduces the amount of memory needed to compute those patterns. The new ASL^{VH} class is also useful for showing how ARs compare with string representations because it crosscuts the subregular stringset hierarchy: ASL^{VH} contains patterns like Akan, Baiyinna Orochen, and Kinande which are SL over strings of vowels; Finnish which is SP over strings of vowels; and Tutrugbu which is LT over strings of vowels. The ASL^{VH} class also excludes the unattested First-Last Harmony (FLH) pattern which is Star Free (SF) and its description requires First Order (FO) logic over strings of vowels.

In addition this dissertation includes one example of an attested vowel harmony pattern which does not fit into the ASL^{VH} class. Eastern Meadow Mari (EMM) could not be described locally due to its use of a harmonic contrast between the two transparent vowels. This contrast makes it impossible to distinguish one of the transparent vowels' back features from the back feature of the harmonizing vowel. Thus a longer string of alternating transparent vowels would result in more iterations of features on the back tier which intervene between the back features associated to the harmonizing vowels. The logical possibility of alternating transparent vowels in an EMM vowel string prevents the pattern from being described locally. The problem that this counterexample presents to the ASL^{VH} hypothesis could have one of two different explanations:

- (a) The ASL^{VH} hypothesis is too restrictive to include all attested vowel harmony patterns and a more inclusive class of ARsets which requires a higher level of logic is needed.
- (b) The multi-tiered ARs I have defined are too restrictive and must be altered or expanded in some way to describe a pattern with a harmonic contrast between transparent vowels in a local way.

A restrictive theory of vowel harmony posits a class of ARsets which includes patterns that are attested in natural language and excludes unattested patterns. Enriching the representation as suggested in (b) is generally more restrictive than increasing the expressivity of a grammar as suggested in (a). Each class that is higher up on the subregular stringset hierarchy is mathematically defined based on the level of logic needed to describe the patterns within it. Each level of logic vastly increases the complexity of the patterns that can be described.

For example, the SF class requires FO logic and includes patterns which phonological theory has deemed pathological because we do not see examples of such patterns in natural language, such as First-Last Harmony (FLH) (Heinz 2018; Lai 2015). A class of ARsets that is more expressive than ASL^{VH} and requires a higher level of logic than a CNL to describe those patterns could also overgenerate to include pathological patterns like FLH. On the other hand, abstract representations can be tweaked and their individual assumptions and properties can be tested. In this way we can determine if ARs can be expanded to include attested vowel harmony patterns like EMM and restricted to exclude unattested patterns like FLH. To solve the problem posed by EMM both explanations (a) and (b) would require additional research into the possible uses and extensions of multi-tiered ARs which is beyond the scope of this dissertation.

While phonological theory has proposed a variety of AR structures this dissertation provides the first theory-independent example of how to study their computational properties. Future work could apply these same FLT methods to a study of different types of autosegmental data structures, such as feature classes (Padgett 1995, 2002) or feature geometry (Clements 1985; Halle 1995; Kornai 1993; McCarthy 1988). Phonological theory could also benefit from a deeper investigation into the effects of representational assumptions on the expressivity of multi-tiered ARs. For example, removing the requirement for full specification; future work could determine whether or not underspecification might provide the necessary expressivity to allow EMM to be ASL^{VH} while still excluding FLH.

Additional future work could include using my multi-tiered ARs to study vowel harmony as an input-output mapping process. If attested surface patterns like EMM cannot be described locally over multi-tiered ARs they might be better understood as an input-output mapping process. Such processes are generally represented as a function which transforms some input structure into a different output structure of the same type. In this dissertation I have shown that vowel harmony patterns without transparent vowel contrasts are local over multi-tiered ARs and their locality is independent of any potential transformational analysis. As a fact of vowel harmony, this restricted notion of locality will strengthen any input-output mapping theory.

This dissertation has shown that representation still matters for phonological theory. Using multi-tiered ARs instead of strings allows one to describe a variety of vowel harmony patterns in a concise local way. In addition, the multi-tiered ARs highlight a specific restriction on autosegmental locality which deserves further investigation. This dissertation provides an explicit formal methodology for studying abstract representational data structures and their computational properties. These methods can help phonologists to better understand the nature of locality in vowel harmony and the usefulness of abstract representations in general.

Bibliography

- Akinlabi, Akinbiyi. 2009. "Neutral Vowels in Lokaa Harmony." *Canadian Journal of Linguistics* 59 (2): 197–228.
- Aksënova, Alena. 2017. "Tier Alphabet(s): Tendencies and Restrictions." Qualifying paper.
- Aksënova, Alena, and Sanket Deshmukh. 2018. "Formal Restrictions on Multiple Tiers." In *Proceedings of the Society for Computation in Linguistics*, 1:Article 8.
- Archangeli, Diana, and Douglas Pulleyblank. 1994. *Grounded Phonology*. Vol. 25. MIT Press.
- . 2002. "Kinande Vowel Harmony: Domains, Grounded Conditions, and One-Sided Alignment." *Phonology* 19 (2): 139–88.
- Baković, Eric. 2000. "Harmony, Dominance, and Control." PhD thesis, New Brunswick, New Jersey: Rutgers University.
- Blum, Eileen. 2019. "On the Locality of Vowel Harmony over Multi-Tiered Autosegmental Representations." Qualifying Paper. Rutgers University.
- Casali, Roderic. 2012. "[+ATR] Dominance in Akan." *Journal of West African Linguistics* 39 (1): 33–59.
- Chandlee, Jane. 2014. "Strictly Local Phonological Processes." PhD thesis, University of Delaware.
- Chandlee, Jane, and Jeffrey Heinz. 2018. "Strict Locality and Phonological Maps." *Linguistic Inquiry* 49 (1): 23–60.
- Chomsky, Noam, and Morris Halle. 1965. "Some Controversial Questions in Phonological Theory." *Journal of Linguistics* 1 (2): 97–138.
- . 1968. *The Sound Pattern of English*. New York: Harper & Row.
- Clements, George N. 1976. "Vowel Harmony in Non-Linear Generative Phonology: An Autosegmental Model." PhD thesis, Harvard University; Bloomington, Indiana University Linguistics Club.
- . 1985. "The Geometry of Phonological Features." *Phonology Yearbook* 2 (1): 225–52.
- . 1986. "Compensatory Lengthening and Consonant Gemination in LuGanda." In *Studies in Compensatory Lengthening*, edited by Leo Wetzels and Engin Sezer, 37–77. Dordrecht: Foris.
- . 1991. "Place of Articulation in Consonants and Vowels: A Unified Theory." *Working Papers of the Cornell Phonetics Laboratory* 5: 77–123.
- Clements, George N., and Samuel Keyser. 1983. *CV Phonology: A Generative Theory of the Syllable*. Cambridge, MA: MIT Press.
- Cole, Jennifer, and Charles Kisseberth. 1994. "An Optimal Domains Theory of Harmony." *Studies in the Linguistic Sciences* 24 (2): 1–13.
- Dolphyne, Florence. 1988. *The Akan (Twi-Fante) Language*. Ghana Universities Press.
- Essegbey, James. 2019. *Tutrugbu (Nyangbo) Language and Culture*. Brill.
- Fidelholtz, James. 1971. "On the Indeterminacy of the Representation of Vowel Length." *Papers in Linguistics* 4: 577–94.
- Gick, Bryan, Douglas Pulleyblank, Fiona Campbell, and Nguessimo Mutaka. 2006. "Low Vowels and Transparency in Kinande Vowel Harmony." *Phonology*, 1–20.
- Goldsmith, John. 1976. "Autosegmental Phonology." PhD thesis, Massachusetts Institute of Technology.
- Halle, Morris. 1995. "Feature Geometry and Feature Spreading." *Linguistic Inquiry* 26 (1): 1–46.
- Hayes, Bruce. 1989. "Compensatory Lengthening in Moraic Phonology." *Linguistic Inquiry* 20: 253–306.
- . 1990. "Diphthongization and Coindexing." *Phonology* 7 (1): 31–71.
- Heinz, Jeffrey. 2010. "Learning Long-Distance Phonotactics." *Linguistic Inquiry* 41 (4): 623–61.
- . 2011a. "Computational Phonology-Part i: Foundations." *Language and Linguistics Compass* 5 (4).
- . 2011b. "Computational Phonology-Part II: Grammars, Learning, and the Future." *Language and Linguistics Compass* 5 (4).
- . 2018. "The Computational Nature of Phonological Generalizations." In *Phonological Typology, Phonetics and Phonology*, edited by Larry Hyman and Frans Plank, 23:126–95. Phonological Typology, Phonetics and Phonology. Mouton de Gruyter.
- Heinz, Jeffrey, and William Idsardi. 2013. "What Complexity Differences Reveal about Domains in Language." *Topics in Cognitive Science* 5 (1): 111–31.
- Heinz, Jeffrey, and Regine Lai. 2013. "Vowel Harmony and Subsequentiality." In *Proceedings of the 13th Meeting on the Mathematics of Language (MoL 13)*, 52–63. Association for Computational Linguistics.

- Heinz, Jeffrey, Chetan Rawal, and Herbert G. Tanner. 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.
- Hyman, Larry. 1985. *A Theory of Phonological Weight*. Dordrecht: Foris.
- . 2002. "Is There a Right to Left Bias in Vowel Harmony?" In *9th International Phonology Meeting*. Vol. 1. Vienna, Austria.
- Ingria, Robert. 1980. "Compensatory Lengthening as a Metrical Phenomenon." *Linguistic Inquiry* 11: 465–95.
- Jardine, Adam. 2016a. "Computationally, Tone Is Different." *Phonology* 33: 247–83.
- . 2016b. "Locality and Non-Linear Representations in Tonal Phonology." PhD thesis, University of Delaware.
- . 2017a. "On the Logical Complexity of Autosegmental Representations." In *15th Meeting on the Mathematics of Language*, 22–35.
- . 2017b. "The Local Nature of Tone Association Patterns." *Phonology* 34 (2): 385–405.
- . 2019. "The Expressivity of Autosegmental Grammars." *Journal of Logic, Language, and Information* 28 (1): 9–54.
- Jardine, Adam, and Jeffrey Heinz. 2015. "A Concatenation Operation to Derive Autosegmental Graphs." In *Proceedings of the 14th Annual Meeting on the Mathematics of Language (MoL 2015)*, 139–51. Chicago, USA: Association for Computational Linguistics.
- Jurgec, Peter. 2011. "Feature Spreading 2.0: A Unified Theory of Assimilation." PhD thesis, University of Tromsø.
- Kaplan, Ronald, and Martin Kay. 1994. "Regular Models of Phonological Rule Systems." *Computational Linguistics* 20 (3).
- Kenstowicz, Michael. 1970. "On the Notation of Vowel Length in Lithuanian." *Papers in Linguistics* 6: 73–113.
- Keyser, Samuel, and Paul Kiparsky. 1984. "Syllable Structure in Finnish Phonology." In *Language and Sound Structure*, edited by Mark Aronoff and Richard Oehrle, 7–31. Cambridge, MA: MIT Press.
- Kornai, Andras. 1993. "The Generative Power of Feature Geometry." *Annals of Mathematics and Artificial Intelligence* 8 (1): 37–46.
- Krämer, Martin. 2001. "Vowel Harmony and Correspondence Theory." PhD thesis, Düsseldorf: Heinrich Heine Universität.
- Lai, Regine. 2015. "Learnable Vs. Unlearnable Harmony Patterns." *Linguistic Inquiry* 46 (3): 425–51.
- Lamont, Andrew. 2019. "Sour Grapes Is Phonotactically Complex." In *LSA 2019 Annual Meeting*.
- Leben, William. 1973. "Suprasegmental Phonology." PhD thesis, Massachusetts Institute of Technology.
- Levergood, Barbara. 1984. "Rule Governed Vowel Harmony and the Strict Cycle." In *Proceedings of NELS*, 14:275–93. University of Massachusetts, Amherst; GLSA.
- Levin, Juliette. 1985. "A Metrical Theory of Syllabicity." PhD thesis, Massachusetts Institute of Technology.
- Li, Bing. 1996. *Tungusic Vowel Harmony: Description and Analysis*. University of Amsterdam; Holland Institute of Generative Linguistics.
- McCarthy, John. 1979. "Formal Problems in Semitic Phonology and Morphology." PhD thesis, Massachusetts Institute of Technology.
- . 1982. "Prosodic Templates, Morphemic Templates, and Morphemic Tiers." In, edited by Harry van der Hulst and Norval Smith, part 1:191–223. Dordrecht: Foris.
- . 1988. "Feature Geometry and Dependency: A Review." *Phonetica* 38. http://scholarworks.umass.edu/linguist_faculty_pubs/38.
- . 2004. "Headed Spans and Autosegmental Spreading." University of Massachusetts Amherst.
- McCarthy, John, and Alan Prince. 1986. "Prosodic Morphology." University of Massachusetts, Amherst; Brandeis University.
- McCawley, James. 1968. *The Phonological Component of a Grammar of Japanese*. The Hague & Paris: Mouton de Gruyter.
- McCollum, Adam, Eric Baković, Anna Mai, and Eric Meinhardt. 2020. "Unbounded Circumambient Patterns in Segmental Phonology." *Phonology* 37 (2): 215–55.
- McCollum, Adam, and James Essegbey. 2020. "Initial Prominence and Progressive Vowel Harmony in Tʉtrugbu." *Phonological Data and Analysis* 2 (3): 1–37.
- McMullin, Kevin. 2016. "Tier-Based Locality in Long-Distance Phonotactics: Learnability and Typology."

- PhD thesis, University of British Columbia.
- McMullin, Kevin, and Gunnar Hansson. 2014. "Long-Distance Phonotactics as Tier-Based Strictly 2-Local Languages." In *Proceedings of the 2014 Annual Meeting on Phonology*.
- McNaughton, Robert, and Seymour Papert. 1971. "Counter-Free Automata." In MIT Research Monographs 65. Cambridge, Mass., MIT Press.
- Mutaka, Ngessimo. 1995. "Vowel Harmony in Kinande." *Journal of West African Languages* 25 (2): 41–55.
- Nevins, Andrew. 2010. *Locality in Vowel Harmony*. *Linguistic Inquiry Monographs*. Vol. 55. MIT Press.
- NiChiosain, Maire, and Jaye Padgett. 1993. "Inherent Vplace." Linguistics Research Center.
- O'Hara, Charlie, and Caitlin Smith. 2019. "Computational Complexity and Sour-Grapes-Like Patterns." In *Proceedings of the 2018 Annual Meeting on Phonology*. Vol. 6.
- O'Keefe, Michael. 2004. "Akan Vowel Harmony." Swarthmore College.
- Odden, David. 1987. "Predicting Tone in Kikuria." In *Current Approaches to African Linguistics*, edited by David Odden, 4:311–26. Dordrecht: Foris.
- . 1994. "Adjacency Parameters in Phonology." *Language* 70 (2): 289–330.
- . 1996. *The Phonology and Morphology of Kimatuumbi*. Oxford: Clarendon Press.
- . 2011. "The Representation of Vowel Length." In *The Blackwell Companion to Phonology*, 1–26. Wiley Online Library.
- Padgett, Jaye. 1995. "Feature Classes." In *Papers in Optimality Theory*, edited by J. Beckman, S. Urbanczyk, and L. Walsh, 18:385–420. University of Massachusetts Occasional Papers.
- . 2002. "Feature Classes in Phonology." *Language* 78 (1): 81–110. <http://www.jstor.org/stable/3086646>.
- . 2011. "Consonant-Vowel Place Feature Interactions." In *The Blackwell Companion to Phonology*, edited by M. Van Oostendorp, C. J. Ewen, E. V. Hume, and K. Rice, 3:1761–86. Wiley Online Library.
- Prince, Alan, and Paul Smolensky. 1993. "Optimality Theory: Constraint Interaction in Generative Grammar." Rutgers University Center for Cognitive Science; Rutgers Optimality Archive.
- Pulleyblank, Douglas. 1996. "Neutral Vowels in Optimality Theory: A Comparison of Yoruba and Wolof." *Canadian Journal of Linguistics* 41: 295–347.
- Pyle, Charles. 1971. "On the Treatment of Length in Generative Phonology." PhD thesis, University of Illinois.
- Ringen, Catherine. 1975. "Vowel Harmony: Theoretical Implications." PhD thesis, Indiana University.
- Ringen, Catherine, and Orvokki Heinamaki. 1999. "Variation in Finnish Vowel Harmony: An OT Account." *Natural Language and Linguistic Theory* 17: 303–37.
- Ringen, Catherine, and Robert Vago. 1998. "Hungarian Vowel Harmony in Optimality Theory." *Phonology* 15: 393–416.
- Rogers, James, Jeffrey Heinz, Gil Bailey, Matt Edlefsen, Molly Visscher, David Wellcome, and Sean Wibel. 2010. "On Languages Piecewise Testable in the Strict Sense." In *The Mathematics of Language*, edited by Christian Ebert, Gerhard Jäger, and Jens Michaelis, 6149:255–65. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg.
- Rogers, James, Jeffrey Heinz, Margaret Fero, Jeremy Hurst, Dakota Lambert, and Sean Wibel. 2013. "Cognitive and Sub-Regular Complexity." *Formal Grammar*, 90–108.
- Rogers, James, and Geoffrey Pullum. 2011. "Aural Pattern Recognition Experiments and the Subregular Hierarchy." *Journal of Logic, Language, and Information* 20: 329–42.
- Sagey, Elizabeth. 1986. "The Representation of Features and Relations in Non-Linear Phonology." PhD thesis, Massachusetts Institute of Technology.
- Smith, Caitlin, and Charlie O'Hara. 2019. "Formal Characterizations of True and False Sour Grapes." In *Proceedings of the Society for Computation in Linguistics*. Vol. 2. 41.
- Sulkala, Helena, and Merja Karjalainen. 1992. *Finnish*. Descriptive Grammars. Routledge.
- Välilmaa-Blum, Riitta. 1986. "Finnish Vowel Harmony as a Prescriptive and Descriptive Rule: An Autosegmental Account." In *Proceedings of the Third Eastern States Conference on Linguistics*, edited by Fred Marshall. University of Pittsburgh.
- van der Hulst, Harry. 2016. "Vowel Harmony." In *Oxford Research Encyclopedia of Linguistics*. Oxford University Press.
- . 2017. "A Representational Account of Vowel Harmony in Terms of Variable Elements and Licensing." In *Approaches to Hungarian*. Vol. 15. Leiden Conference. John Benjamins Publishing Company.
- . 2018. "Opacity and Transparency in Vowel Harmony." In *Asymmetries in Vowel Harmony: A*

- Representational Account*, 3–50. Oxford University Press.
- van der Hulst, Harry, and Norval Smith. 1986. “On Neutral Vowels.” In *The Phonological Representation of Suprasegmentals: Studies on African Languages*, 4:233–81.
- van der Hulst, Harry, and Jeroen van de Weijer. 1995. “Vowel Harmony.” In *The Handbook of Phonological Theory*, edited by John A. Goldsmith, 495–534. Blackwell.
- Vaysman, Olga. 2009. “Segmental Alternations and Metrical Theory.” PhD thesis, Massachusetts Institute of Technology.
- Walker, Rachel. 2010. “Nonmyopic Harmony and the Nature of Derivations.” *Linguistic Inquiry* 41 (1): 169–79.
- . 2011. *Vowel Patterns in Language*. Vol. 130. Cambridge University Press.
- . 2014a. “Nonlocal Trigger-Target Relations.” *Linguistic Inquiry* 45 (3): 501–23.
- . 2014b. “Surface Correspondence and Discrete Harmony Triggers.” In *Proceedings of the Annual Meeting on Phonology*.
- White, Robin. 1972. “Klamath Phonology.” PhD thesis, University of Washington.
- Zec, Draga. 1988. “Sonority Constraints on Prosodic Structure.” PhD thesis, Stanford University.