

Input and output locality and representation

Jane Chandlee* and Adam Jardine†

DRAFT
April 1, 2020

Abstract

Using a rigorous, computational notion of locality, this paper evaluates one of the central motivations for autosegmental representations (ARs)—that they reduce long-distance processes to local ones. We analyze a variety of tonal processes, both in isolation and in their interactions, using the *autosegmental input-strictly (A-ISL) local functions* of Chandlee and Jardine (2019). This reveals that ARs do describe long-distance processes locally in some cases. We further show that these processes are *necessarily* input-local, meaning that they cannot be captured with an output-based notion of locality. However, there are cases when changes are based on output information that are not A-ISL. We further discuss how the A-ISL class may be integrated into a fuller computational theory of phonology that also includes output locality. Altogether then, a precise computational analysis of ARs provides crucial insights into the interactions among tone, locality, and phonological computation.

1 Introduction

It has long been recognized that locality plays a central role in phonology, as various notions of what it means to be local have been proposed as restrictions on possible grammars. Preserving locality was in fact one of the major motivating factors behind the use of feature-geometric, autosegmental representations (ARs; Williams, 1976; Goldsmith, 1976; Clements, 1977). As Odden (1994, pg. 289) states, the advantage of ARs is that they “[make] it possible to view apparently long-distance rules as rules operating between segments which are adjacent at a specified level, even though the segments are not adjacent at all levels of representation”. However, since the advent of Optimality Theory (OT; Prince and Smolensky, 1993), the utility of nonlinear representations has explicitly come into question (Leben, 2006; Hyman, 2014; Shih and Inkelas, 2019).

At the same time, recent work in theoretical computational linguistics has made it possible to rigorously study representations and locality. Theoretical computer science offers precise definitions of the term “local” based on principles of computation that can be applied to natural language (Rogers and Pullum, 2011; Chandlee and Heinz, 2018). Briefly, a local computation is one which operates over some sequence of k positions in a string, as exemplified in (1). For example, a *CLASH constraint against adjacent H-toned syllables (Zoll, 2003), as in (2a), is local in this sense because it evaluates sequences of length 2.

(1) $\boxed{a_1} \boxed{a_2} \dots \boxed{\tau_1} \boxed{\tau_2} \dots \boxed{\tau_k} \dots \boxed{a_n}$

- (2) a. * $\acute{\sigma} \acute{\sigma}$
b. $\sigma \sigma \sigma \acute{\sigma} \acute{\sigma} \sigma$
c. $\sigma \acute{\sigma} \sigma \sigma \acute{\sigma} \sigma$


The string in (2b) violates (2a) because it contains the forbidden 2-sequence, whereas (2c) does not.

A long-distance Obligatory Contour Principle-style constraint which bans two H-tones in a particular domain (i.e., that both (2b) and (2c) violate), however, is not local in this conception—at least over surface strings. For phonotactics, recent work has shown the utility of tier projections (Heinz et al., 2011; McMullin and Hansson, 2016) and ARs (Jardine, 2017a, 2019) in capturing long-distance patterns with grammars that are local in this way. For example, the constraint in (2a) can be reformulated over ARs as a constraint

*Haverford College, jchandlee@haverford.edu

†Rutgers University, adam.jardine@rutgers.edu

over two H tones adjacent on the tonal tier; this is violated both by adjacent H-toned tone bearing units (TBUs), as in (3b), or H tones that are associated to non-adjacent TBUs (3c). This is because both contain the forbidden HH sequence *on the tonal tier*. In this way, ARs capture long-distance phonotactic patterns with local constraints.



This conception of locality has been applied to tone processes as well. Chandlee and Jardine (2019a) propose the *autosegmental input-strictly local* (A-ISL) class of functions that transform underlying ARs to surface ARs, based on Chandlee (2014)’s *input strictly local* (ISL) functions for strings. These classes formalize a notion of locality in which any change is based on information in the input that is within some fixed distance of that change. Thus, the computational perspective allows for a rigorous study of ARs and to what extent they make phonological processes local. That is, to determine if ARs do contribute to a theory of phonology that is fundamentally local, we can examine non-ISL (i.e., non-local over strings) phonological processes that are A-ISL (i.e., local over ARs).

The purpose of this paper is to demonstrate, from this formal perspective, *how* ARs work to preserve phonological locality. Studying locality from such a perspective will reveal that 1) perhaps not surprisingly, ARs make some, but not all, long-distance processes local, and 2) tonal phonology requires not just ARs, but both input- and output-based notions of locality. In particular, we argue that the A-ISL function class is a *necessary*, but not *sufficient*, local characterization of tone. That is, there are tone processes that are A-ISL but not ISL, meaning ARs do enable a local analysis of certain non-local phenomena. As has been argued in previous literature on ARs, this draws from the fact that long-distance changes can be described by local manipulations of elements on distinct, asynchronous tiers and the association relation between them.

However, there are also ISL processes that are *not* A-ISL, meaning ARs can actually *prevent* a local analysis. These are cases in which a change on one tier depends on information in another. Finally, there are also intuitively “local” processes that are in fact neither ISL nor A-ISL because they are based on output, and not input, locality. These subtle but important distinctions are only made clear using a precise definition of locality such as the one provided by the ISL and A-ISL function classes. And as we discuss, the existence of processes that are properly A-ISL (that is, A-ISL only) indicates that ARs do indeed make a non-trivial contribution to a local theory of phonology. We further discuss how the A-ISL class may be integrated into a fuller computational theory of phonology that also includes *output* locality.

1.1 Overview of the proposal

To briefly introduce how A-ISL can capture non-local processes (with a more thorough explanation to follow in the next section), consider unbounded shift in Zigula (Kenstowicz and Kisseberth, 1990), in which an underlying H tone shifts to the penultimate TBU in the word.

- (4) Zigula (Kenstowicz and Kisseberth, 1990)
- | | | | |
|----|---------------------|--------------------|-------------------------|
| a. | /ku-songoloz-a/ | [ku-songoloz-a] | ‘to avoid’ |
| b. | /á-songoloz-a/ | [a-songolóz-a] | ‘He/she is avoiding’ |
| c. | /ku-lómbez-a/ | [ku-lombéz-a] | ‘to ask’ |
| d. | /ku-lómbez-ez-an-a/ | [ku-lombezez-án-a] | ‘to ask for each other’ |

This is a non-local process because whether or not a penultimate TBU surfaces as H-toned depends on whether or not there is a preceding TBU bearing an H tone in the underlying form. Thus, the process is not ISL (as we show below). However, it is A-ISL. As an AR process, we can view this shift as the reassociation of an H tone to the penultimate TBU in the word, as shown below in (5) for (4b). (Here, as in the rest of the paper, without loss of generality, we use syllables as a representative TBU.)

$$(5) \quad \begin{array}{ccc} \# \text{ H } \# & \mapsto & \# \text{ H } \# \\ | & \dashrightarrow & \\ \# \sigma \sigma \sigma \sigma \sigma \# & & \# \sigma \sigma \sigma \sigma \sigma \# \end{array}$$

Following Chandlee and Jardine (2019a), we define A-ISL using *quantifier-free (QF) logical transductions* (Chandlee and Lindell, forthcoming). QF transductions reference input representations using a predicate logic in which variables range over elements (or *positions*) in the representation—but quantifiers are disallowed. For example, we can define the output association relation $A'(x, y)$ in Zigula as follows.

$$(6) \quad A'(x, y) \stackrel{\text{def}}{=} \text{H}(x) \wedge \sigma(y) \wedge \#(s(x)) \wedge \#(s(s(y)))$$

$$(7) \quad \begin{array}{ccc} \# \boxed{\text{H}} \# & \mapsto & \# \boxed{\text{H}} \# \\ | & \dashrightarrow & \\ \# \sigma \sigma \sigma \boxed{\sigma} \sigma \# & & \# \sigma \sigma \sigma \boxed{\sigma} \sigma \# \end{array}$$

We will explain the details of logical transductions in the next section, but essentially (6) says that an output TBU y should be associated to a high tone x if and only if it is two positions away from the end of the word ($\#(s(s(y)))$), where s is the *successor* function that establishes the linear order of positions. More generally, formulas such as these establish the output form in terms of information in the input, and that information is limited to a fixed window around particular input positions (in our case TBUs and tones). These input windows are highlighted in (7). The significance of the formula being QF is that global information that would require scanning the entire input (e.g., ‘there exists a position such that...’) is not available. In this way the transduction is required to operate in a local manner.

A consequence of the definition of A-ISL is that locality is also enforced in terms of how the two tiers of an AR can interact. In particular, 1) for any input element x (on either tier), its output value depends on local input information on its same tier; and 2) for any input x and y whether or not they are associated is dependent on input information local to x on its tier and y on its tier. As Chandlee and Jardine (2019a) show, this means that for any A-ISL map, the changes on each tier are themselves ISL maps. However, as we show below, A-ISL can capture processes that are not ISL as string maps, such as unbounded tone shift and unbounded tone deletion (a variant of Meussen’s rule). This confirms, from a formal perspective, that the power of ARs comes from the asynchronicity of distinct tiers and the manipulation of the association relation between them. Thus, while the A-ISL class is just one possible way of defining a notion of locality over ARs, A-ISL minimally illustrates the traditional conception of ARs as structures that capture non-local processes locally. In addition to analyzing individual tone processes to make this point, we also provide an analysis of multiple interacting processes in Shona to show how logical definitions of the A-ISL class can capture multiple generalizations in a unified way.

One example of an alternative notion of locality is that of *output* locality—that a change in a position is based on a local window in the *output*. This is, for example, the notion of locality formalized by Markedness constraints in OT. A class that incorporates output locality in a parallel way to the ISL class is *output strictly local*, or OSL, functions (Chandlee et al., 2015a). The OSL functions are those in which any change is based on information in the output that is within some fixed distance from the change. For example, *iterative* local processes, such as unbounded tone spread, are OSL and not ISL. As we discuss below, unbounded spread is thus also not A-ISL. However, we do not take this as evidence that input locality is the wrong notion of locality and should be abandoned for output locality, as certain processes (i.e., bounded spread and bounded shift) are *necessarily* input-based (as revealed by Myers 1997’s OT analysis of Rimi, which we discuss below). Also, our analysis of unbounded shift, which is A-ISL but not OSL, shows that for some processes the solution to preserving locality is not to consider output locality, but to consider locality over ARs. Thus, while this paper will not offer a complete theory of tone processes, it will point out in A-ISL a class of functions that *must* be included in such a theory. Toward the end we discuss some possibilities for defining classes that include both A-ISL and output-local functions over ARs.

1.2 Overview of the paper

The paper is organized as follows. §2 defines the computational property of input strict locality (ISL), explains how it is used in a logical transduction of a phonological process, and extends it to processes that operate over ARs (i.e., processes that are A-ISL). In §3 we apply these notions of locality to analyses of several types of tone processes, including bounded and unbounded tone shift and tone spread in §3.1 and variants of Meussen’s rule in §3.2. §4 then gives a more in-depth analysis of Meussen’s in Shona, which

interacts in various ways with the OCP. The result of this focused analysis on a single language is that a logical definition of association can capture a cluster of related local processes in a unified way. In §5 we discuss how ISL and A-ISL compare to previous notions of locality in phonology, and also discuss how to incorporate output locality into the proposal. §6 concludes.

2 Input strict locality

The ISL functions are those that determine an output string for a given input string based only on input substrings of bounded length (Chandlee, 2014; Chandlee et al., 2014). As a subset of the regular relations, they have a reduced computational complexity and expressivity. Despite their restrictive nature, however, they have been shown to be sufficient to model a significant range of segmental processes. The property of ISL thus serves as a precise, computationally defined notion of phonological locality. Here we provide a brief, informal explanation of what it means to be ISL; readers are referred to the work cited above for the technical details and for formal language-theoretic and automata-theoretic characterizations.

The defining trait of an ISL map is essentially this. Given an input string, each segment of that input string contributes some (possibly empty) string to the output string, and what that output string is is locally-determined. In other words, each input segment’s contribution to the output string is based only on the segment itself and a bounded number of its surrounding segments. This is illustrated in the example place assimilation map shown in Figure 1. Given an input like /ɪmpɹɪəbəl/, the function determines the corresponding output string by only focusing on a bounded substring in the input at any given time during the computation. Thus, as highlighted in Figure 1, when it reads the /np/ sequence in the input, it can append [mp] to the output. It doesn’t matter what came before this /np/ and it doesn’t matter what might come after it: the output at this stage must always include [mp]. ISL functions are parameterized by the length of the substring it needs to look for; in this example the length of that substring (/np/) is 2, so the function is a 2-ISL function.

Input:	ɪ	n	p	ɹ	ɪ	ə	b	ɹ	b	l
Output:	ɪ		mp	ɹ	ɪ	ə	b	ɹ	b	l

Figure 1: ISL nasal place assimilation map for the input ‘improbable’

The idea that local phonology pays attention to contiguous substrings is certainly not novel, but ISL provides a precise notion of what it means for a phonological map to be local in a computational sense. Our goal in what follows is to see how well the notion of locality provided by ISL aligns with the sense in which autosegmental representations enable a local treatment of otherwise non-local phenomena. In pursuit of this goal we will use the logical characterization of ISL, because it enables a more direct comparison among different representations (in our case strings versus ARs).

To demonstrate the logical characterization of ISL, first over strings, we will use the example of bounded tone shift, in which a tone shifts some fixed number of TBUs away from its underlying position. For example in Rimi, a tone shifts one TBU to the right. In (8), a suffix surfaces with a high tone when attached to a stem that ends in a high tone, and the stem itself no longer bears that tone.

- (8) Rimi (Schadeberg, 1979; Myers, 1997)
- a. /u-hang-a/ [u-hang-a] ‘to meet’
 - b. /u-púm-a/ [u-púm-á] ‘to go away’
 - c. /mu-ntu/ [mu-ntu] ‘person’
 - d. /rá-mu-ntu/ [ra-mú-ntu] ‘of a person’
 - e. /u-huvi-ì/ [u-huvi-ì] ‘belief’
 - f. /mu-tém-ì/ [mu-tem-ì] ‘chief’

In line with the assumption that tone processes operate over TBUs, we model this case of tone shift using strings of syllables instead of strings of segments. The examples in (8) can then be rewritten as in (9).¹

- (9)
- a. $\sigma\sigma\sigma \mapsto \sigma\sigma\sigma$
 - b. $\sigma\acute{\sigma}\sigma \mapsto \sigma\sigma\acute{\sigma}$
 - c. $\sigma\sigma \mapsto \sigma\sigma$

¹Throughout the paper we reserve the symbol \rightarrow for rewrite rules and use \mapsto for example input-output maps of a given function.

- d. $\acute{\sigma}\sigma\sigma \mapsto \sigma\acute{\sigma}\sigma$
- e. $\sigma\sigma\sigma \mapsto \sigma\sigma\sigma$
- f. $\sigma\acute{\sigma}\sigma \mapsto \sigma\sigma\acute{\sigma}$

The logical characterization of ISL defines formulas in first order (FO) logic that evaluate to true for elements of the input string of a particular type.² For example, we might use a formula like $\acute{\sigma}(x)$ to pick out input elements that bear a high tone (such as the second element in the input string in (9b)). We can also use terms like $p(x)$, or *predecessor* of x , and $s(x)$, or *successor* of x , to refer to elements immediately preceding or following x , respectively. Lastly, formula marked with prime evaluate to true for elements of the *output* string. For example, the formula in (10) specifies the output element that bears a high tone.

$$(10) \quad \acute{\sigma}'(x) \stackrel{\text{def}}{=} \acute{\sigma}(p(x))$$

(10) says that an output element should bear a high tone if its predecessor in the *input* bears a high tone. To illustrate, (11) highlights in boxes exactly the TBUs from (9a), (9b), and (9d) that satisfy (10) and thus will be labeled as H-toned in the output.

$$(11) \quad \text{a. } \sigma\sigma\sigma \quad \text{b. } \sigma\acute{\sigma}\boxed{\sigma} \quad \text{c. } \acute{\sigma}\boxed{\sigma}\sigma$$

The formula in (12) identifies those output elements that are unspecified for tone: namely, those elements whose predecessors in the input do not bear a high tone.

$$(12) \quad \sigma'(x) \stackrel{\text{def}}{=} \neg\acute{\sigma}(p(x))$$

Notice that the elements in (11) that do not satisfy $\acute{\sigma}'(x)$ do satisfy $\sigma'(x)$; thus, every element receives some label in the output. Following Engelfriet and Hoogetboom (2001), if some element does not receive a label, it is deleted in the output. We make use of this in §3.2.

The use of FO formulas in this way corresponds to an ISL function because the formulas are importantly *quantifier-free* (QF). This limits their ability to identify the conditions under which an alternation can take place: in Rimi determining that an output element bears a high tone was a matter of checking its immediate neighbor in the input. This is shown explicitly in (13), which highlights two TBUs from (9b) and shows how their surface value is determined by the value in the input of the preceding TBU (highlighted in gray).

$$(13) \quad \text{a. } / \sigma \acute{\sigma} \boxed{\sigma} /$$

$$\quad \quad \quad \downarrow$$

$$\quad \quad \quad [\sigma \sigma \boxed{\acute{\sigma}}]$$

$$\text{b. } / \boxed{\sigma} \acute{\sigma} \sigma /$$

$$\quad \quad \quad \downarrow$$

$$\quad \quad \quad [\sigma \boxed{\acute{\sigma}} \acute{\sigma}]$$

The predecessor and successor functions can be embedded, such as $p(p(x))$ to identify the predecessor of a predecessor, but importantly the degree of embedding must always be finite. This ensures that the conditions used to identify output elements exist within a bounded ‘window’ of information in the input. Chandlee and Jardine (2019b) show that these transductions describe exactly the ISL functions, given that the transductions are *order-preserving*, meaning that such that a position that precedes/follows another position in the input will also precede/follow that position in the output. We follow this assumption throughout the paper.

This sense of locality provides a precise means for identifying what are considered ‘long-distance’ processes in phonology. Consider a pattern like the consonant agreement found in Kikongo:

$$(14) \quad \text{Kikongo (Ao, 1991)}$$


$$\quad /tu-nik-idi/ \mapsto [tunikini], \text{ ‘we ground’}$$

The suffix *-idi* surfaces as *-ini* when it attaches to a stem that contains a nasal consonant. As that nasal consonant can be an unbounded number of segments away from the suffix consonant itself, this pattern is considered long-distance. In terms of the FO characterization used in this paper, determining the nasality of the suffix consonant would require a quantifier to examine the entire stem. The use of embedded predecessor functions is not possible as there is no upper bound on how many preceding segments must be examined to confirm the presence or absence of the stem nasal. In this way the ISL property formalizes our intuitive

²See Chandlee and Jardine (2019a) for a more formal treatment.

understanding of what it means for a process to be local versus long-distance.

Given this means of distinguishing local and long-distance, we can examine how long-distance patterns are made local by ARs. We first must extend ISL to operate over ARs instead of strings. To illustrate with Rimi, as an autosegmental rule, bounded shift involves both delinking and reassociation to a following syllable, as in (15a). The AR for the map in (9d) /rá-mu-ntu/ \mapsto [ra-mú-ntu] ‘of a person’ is given in (15b).




Compared to string representations, ARs consist of two strings (one of tones and one of TBUs), each with their own predecessor and successor functions. In addition, an association relation A is defined that links tones to TBUs. Transductions over ARs are also achieved using formula that define output items in terms of the input structure; these formula might refer to the tones, TBUs, or the associations among them. In the case of Rimi, the transduction involves a change in the association relation from input to output, using the formula in (16).

$$(16) \quad A'(x, y) \stackrel{\text{def}}{=} A(x, p(y))$$

This formula says simply that a TBU is associated to a tone in the output if that tone was associated to its predecessor in the input. Again, this formula is QF, which enforces locality in that same way it did for the maps over string representations. That is, the association between x and y is dependent entirely on information local to x and y in the input.

The example in (17) demonstrates how this formula implements the tone shift. The output association relation A' is indicated (as in Fig. 2) with dotted lines. In this example, $A'(x, y)$ is only true when x is interpreted as the H and y as syllable 2, and so these two positions are associated in the output. The formula does not hold when y is interpreted as any other syllable, and so the H is not associated to any other syllables.



Note that a full definition of this function would also include the formulas $H'(x)$ and $\sigma'(x)$, for when an input element surfaces as H and σ , respectively. However, as these do not change—that is, their definitions would be $H(x)$ and $\sigma(x)$, respectively—for the sake of brevity we will omit from all analyses the formulas for any aspects of the representation that do not change from input to output.

A point we should raise here with respect to representation is that the ARs in (17), like those in (15a), assume that, at the phonological level, TBUs are either associated to an H tone or are unspecified. This follows the analysis in Myers (1997). In the same fashion, the computational analyses in this paper follow the representational assumptions of the sources from which the original analyses are taken. This allows us to study the computational properties of the processes as they were originally posited. In many cases, this means assuming an H versus unspecified contrast, as in Rimi, but in other cases this means an H versus L contrast, or even H versus L versus unspecified. The computational implications of adopting different assumptions for specification of TBUs is an interesting, yet separate question that we leave for future work.

In what follows we will survey a set of tone processes using these formal notions of ISL and A-ISL. Specifically, we will classify patterns using the definitions in (18).

- (18) a. A process is **ISL** if it can be modeled as a QF FO logical transduction over string representations.
 b. A process is **A-ISL** if it can be modeled as a QF FO logical transduction over ARs.

The relationship between ISL and A-ISL is further established by the following theorem:

- (19) **Theorem.** If an AR map is A-ISL, then the individual map on each tier is ISL.
(Chandlee and Jardine, 2019a, Thm 1)

In other words, the changes on each tier in an A-ISL map can be described themselves as ISL functions. This means those processes for which a change in tone depends on information on both tiers are not A-ISL (though they may still be ISL). As we will show, such changes cannot be modeled without the use of quantifiers.

Our survey will reveal processes in all possible categories: ISL only, A-ISL only, both ISL and A-ISL, and neither ISL nor A-ISL. The resulting conclusion is that ARs do not automatically render non-local patterns local, which in itself is not controversial. The larger contribution is a more precise metric by which to assess the conditions under which ARs *do* render non-local patterns local.

3 Local and non-local tone processes

We now survey major types of tone patterns and ask the following two questions:

- (20) a. Is the process ISL; that is, is it local given a string representation?
b. Is the process A-ISL; that is, is it local over ARs?

The set of patterns we will analyze include: tone spread and tone shift, both bounded and unbounded (§3.1), and various forms of Meussen’s rule (§3.2). Importantly, we treat each process as an independent function, in order to assess its locality in isolation. In most cases, this means assuming a single underlying H. We will also consider a set of interacting processes in a single language in §4.

3.1 Spread and shift

We first analyze *bounded spread* and *bounded shift*. In the former, as exemplified below with a pattern from Northern Bemba (Bickmore and Kula, 2013), an underlying H tone spreads to the some fixed number of TBUs. In Northern Bemba it spreads exactly one TBU to the right; this is also referred to as ‘binary tone spread’ or ‘tone doubling’ (Bickmore and Kula, 2013).³

- (21) Northern Bemba (Bickmore and Kula, 2013)
- | | | | |
|----|---------------|---------------|--------------------|
| a. | /tu-la-kak-a/ | [tu-la-kak-a] | ‘we tie up’ |
| b. | /bá-la-kak-a/ | [bá-lá-kak-a] | ‘they tie up’ |
| c. | /bá-ka-fik-a/ | [bá-ká-fik-a] | ‘they will arrive’ |
| d. | /bá-ka-bil-a/ | [bá-ká-bil-a] | ‘they will sew’ |

Intuitively, this is a local process, because the target syllable is adjacent to its trigger. We can verify this formally by observing that it is ISL when viewed as a process operating over strings of syllables. A QF formula for this function is given below in (22).

$$(22) \quad \acute{\sigma}'(x) \stackrel{\text{def}}{=} \acute{\sigma}(x) \vee \acute{\sigma}(p(x))$$


As shown below in the example in (23), the above formula assigns an output high tone ($\acute{\sigma}'(x)$) to a syllable that was either high in the input ($\acute{\sigma}(x)$) or to a syllable whose predecessor was high in the input ($\acute{\sigma}(p(x))$).

$$(23) \quad \acute{\sigma}\sigma\sigma\sigma \mapsto \acute{\sigma}\acute{\sigma}\sigma\sigma$$

In the above example, the first syllable is high in the input, so it is also high in the output. Additionally, the predecessor of the second syllable (the first syllable) is high in the input, so the second syllable is also high in the output. None of the other syllables satisfy the formula in (22), so they remain unspecified. Thus, bounded spread is ISL.

It is also local over autosegmental representations. In terms of ARs, bounded spread is implemented by an additional association from an underlying H tone to the TBU to its right, as given in the rule in (24a). The AR for the map in (21b) /bá-la-kak-a/ \mapsto [bá-lá-kak-a] ‘they tie up’ is given in (24b).

³Copperbelt Bemba has ternary spread, which spreads two TBUs to the right.



In logical terms, this bounded spreading rule can be described by the following definition.

$$(25) \quad A'(x, y) \stackrel{\text{def}}{=} \underbrace{A(x, y)}_{(a)} \vee \underbrace{A(x, p(y))}_{(b)}$$

This definition is read as follows: x and y are associated in the output if and only if (a) x and y are associated in the input; or (b) x and the predecessor of y are associated in the input. How this works is illustrated in Fig. 2.




Figure 2: An example of applying the bounded spread formula in (25) to an AR with three syllables. The input association relation A is depicted with a single solid line on the left-hand side; output association A' is depicted with dashed lines on the right-hand side.

The AR in Fig. 2 has three syllables, labeled 1, 2, and 3. Syllable 1 is associated to the H in the output. As such, the H and syllable 1 satisfy $A(x, y)$, disjunct (a) of (25), and are thus associated in the output. Additionally, as the H is associated to the predecessor of syllable 2 (namely, syllable 1), the H and syllable 2 satisfy $A(x, p(y))$, disjunct (b) of (25). This pair is thus also associated in the output. As the H and syllable 3 satisfy neither disjunct, they are not associated in the output.

Bounded spread is thus describable by a QF formula over ARs, and is thus also A-ISL.⁴ It bears repeating that this analysis considers the process in isolation, where in reality, bounded spread in Northern Bemba is blocked by the OCP (Bickmore and Kula, 2013). However, including this constraint would not change the fact that it is A-ISL. For discussion of the interaction of spreading and the OCP, see the discussion in §4 of a nearly identical process in Shona.

As with bounded spread, bounded shift is also A-ISL, as shown in §2 with Rimi. Another example of bounded shift is found in Kuki-Thaadow, in which a string of tones each associate to the following vowel, as in (26). The first and last tones remain associated to their input TBUs.

- (26) Kuki-Thaadow (Hyman, 2011)
- a. /kà zóŋ lién thúm/ \mapsto [kà zòŋ lién thũm] ‘my three big monkeys’
 - b. #òóóó# \mapsto #òòóó#

As a map over strings, Kuki-Thaadow bounded shift can be handled with the set of formulas in (27).


- (27) a. $\delta'(x) \stackrel{\text{def}}{=} (\delta(x) \wedge \#(p(x))) \vee \delta(p(x))$
 b. $\acute{\sigma}'(x) \stackrel{\text{def}}{=} (\acute{\sigma}(x) \wedge \#(p(x))) \vee \acute{\sigma}(p(x))$
 c. $\check{\sigma}'(x) \stackrel{\text{def}}{=} \#(s(x)) \wedge \acute{\sigma}(x) \wedge \delta(p(x))$
 d. $\hat{\sigma}'(x) \stackrel{\text{def}}{=} \#(s(x)) \wedge \delta(x) \wedge \acute{\sigma}(p(x))$

⁴Ternary spread (as in Copperbelt Bemba) is also A-ISL, as witnessed by the formula $A'(x, y) \stackrel{\text{def}}{=} A(x, y) \vee A(x, p(y)) \vee A(x, p(p(y)))$. This is simply the formula in (25) with an additional disjunct $A(x, p(p(y)))$, which is true for an H tone x associated to a syllable two syllables before y . Note that the H and syllable 3 in Fig. 2 satisfy this disjunct, and thus would be associated in the output.

(27a) says that TBUs bear low tones in the output if either they are the first TBU and were low in the input *or* their predecessor was low in the input. This captures both the shift and the fact that the first tones remain the same from input to output. Note the use of # to refer to the first (last) position in the string, which will be the predecessor (successor) of the first (last) TBU.⁵ (We include these boundaries in the diagrams below, and wherever they are crucial to the map, but will not include them in all diagrams.) The parallel formula in (27b) designates those TBUs that bear a high tone in the output.

With a string representation, contour tones must be handled with their own formulas, separate from those for H and L tones. In this case the formula in (27c) says that a TBU bears a LH contour if it is the last TBU, has a high tone in the input, and its predecessor has a low tone. Again a parallel formula (in (27d)) designates output TBUs with an HL contour. And again all of these formulas are QF, and so the Kuki-Thaadow bounded shift map is ISL.

It is also A-ISL. An AR map for the example in (26) is shown in (28), and the formula defining it is given in (29).



(29) $A'(x, y) \stackrel{\text{def}}{=} ((\#(p(y)) \vee \#(s(y))) \wedge A(x, y)) \vee A(x, p(y)))$

Informally, (29) says that a tone x and a TBU y should be associated in the output if y is either the first or last TBU or x is associated to y 's predecessor in the input. Again this formula is QF, and so we can classify the process as A-ISL.

A couple of points are worth noting regarding these analyses of Kuki-Thaadow. First, aside from the disjunct that handles the preservation of tones at the word boundaries, the formulas for bounded shift in both Rimi and Kuki-Thaadow ((16) and (29)) are identical. This is not surprising, as they are two examples of the same process, and yet their example maps (in (17) and (28)) look quite different. This is due in part to the differences in underlying forms—in Rimi a single underlying H tone is assumed while no such restriction holds in Kuki-Thaadow—highlighting another way in which this computational notion of locality interacts with representation. Second, as already stated, our goal for analyzing processes first with strings and then with ARs is to understand the connection between locality and ARs, rather than to assess which analysis (string versus AR) is preferable. But the Kuki-Thaadow analyses in particular reflect another of the motivations for introducing ARs into the theory, which was to simplify the account of certain phenomena. With strings we need the four similar yet distinct formulas in (27), whereas with ARs we just need a single formula to adjust association.

The logical characterization also reveals something about relative complexity when we compare the analyses of bounded shift and spread to their rule-based formulations. Note that the rule for shift in (15a) includes the two operations of spreading plus delinking—in fact, Odden (2001, pg. 76) argues for “a general decomposition of the process of tone shift into tone spread interacting with tonal delinking, rather than including the operation of shift in the formal repertoire of primitive phonological operations”. This suggests that shift is a more complex phenomena than spreading. However, there is no difference in the logical complexity between spread and shift. In fact, the formula for shift is a simplification of that for spread.

This discrepancy between the complexity of the rule formalism and the logical complexity is more stark when considering unbounded spreading and shift. While neither is local in the usual sense—i.e., neither is local over strings—when viewed as AR processes, unbounded shift is A-ISL, whereas unbounded spread is not. We begin with unbounded shift.

In unbounded shift a tone shifts to a particular position in the word (or phrase), no matter how far that position is from the tone's underlying position. One case is in Zigula (Kenstowicz and Kisseberth, 1990), in which a single H tone shifts to the penultimate TBU in the word.

(30) Zigula (Kenstowicz and Kisseberth, 1990)

⁵See Chandlee and Jardine (2019a) for an alternative method for marking word boundaries that uses the predecessor and successor functions instead of explicit boundary symbols.

- | | | | |
|----|---------------------|---------------------|-------------------------|
| a. | /ku-songoloz-a/ | [ku-songoloz-a] | ‘to avoid’ |
| b. | /ku-lómbez-a/ | [ku-lombéʒ-a] | ‘to ask’ |
| c. | /ku-lómbez-ez-a/ | [ku-lombeʒ-éz-a] | ‘to ask for’ |
| d. | /ku-lómbez-ez-an-a/ | [ku-lombeʒ-ez-án-a] | ‘to ask for each other’ |
| e. | /á-songoloz-a/ | [a-songolóʒ-a] | ‘He/she is avoiding’ |


In (30a) the toneless verb root /songoloz/ surfaces with no change, in contrast with the H-toned verb root (30b) /lómbez/ in which the H-tone surfaces on the penultimate syllable.⁶ The examples (30c) and (30d) show that the H tone of the verb truly is shifting to the penult, regardless of whether this is part of the verb root. Further evidence is given in (30e), which shows that the addition of a toned prefix /á/ to a toneless root results in a surface high tone on the verb root, as that is where the penult lies.

Intuitively, this is a long-distance process, because the H-tone shifts unboundedly far away from its initial position. This bears out formally: over strings, this is not an ISL function. In logical terms, this means there is no QF statement $\acute{\sigma}'(x)$ that can accurately determine when an input syllable should appear as high in the output. To see why, consider the contrast between the penultimate syllable in an underlyingly toneless word (31a) and the penult in underlyingly toned words (31b).


- (31) a. $\sigma\sigma\sigma\sigma \mapsto \sigma\sigma\sigma\sigma$
 b. $\sigma\acute{\sigma}\acute{\sigma}\sigma \mapsto \sigma\sigma\acute{\sigma}\sigma$
 $\acute{\sigma}\sigma\sigma\sigma \mapsto \sigma\sigma\acute{\sigma}\sigma$

The formula $\acute{\sigma}'(x)$ for determining an output H-toned syllable should be false for the penult in (31a) but true for the penults in both strings in (31b). In terms of the shifting process, this is because no H-toned syllable precedes the penult in (31a), but one does in both cases in (31b). Thus, $\acute{\sigma}'(x)$ should be true if and only if x is the penult *and* some H-toned syllable precedes x in the input. However, there is no QF formula that can capture this. Intuitively, this is because there is no finite sequence $\acute{\sigma}(p(x)) \vee \acute{\sigma}(p(p(x))) \vee \acute{\sigma}(p(p(p(x)))) \vee \dots$ that will capture all cases in which an H tone syllable precedes x at *some point* in the input. Thus, the unbounded nature of the pattern means as a string function it is not QF-definable and thus not ISL.

Viewed in terms of ARs, unbounded shift has been derived multiple ways. The first, and the one taken by Kenstowicz and Kisseberth (1990), is similar to the accentual derivation described in §5.1 for Ndebele: an accent (or the head of a trochaic foot) is assigned to the penultimate syllable, and the H is reassigned to that syllable.



Another, argued for by Odden (2001), is the decomposition of shift into iterative, unbounded spread and delinking of all but the rightmost association, as shown by the rules and example derivation below.



Evidence for this view can be found in languages in which intervening depressor consonants (usually voiced obstruents) affect the delinking process but not the spreading process (as in, e.g., Digo; Kisseberth, 1984).

⁶For evidence that the H is underlyingly root-initial, see Kenstowicz and Kisseberth (1990).

Regardless, when viewed as a map directly from input to output, the function is one in which the final H associates to the penultimate syllable. This *is* QF over ARs, and thus A-ISL. First, we need to define the concept of a final element in a string, and from that, the notion of a penultimate element.

- (34) a. $\text{final}(x) \stackrel{\text{def}}{=} \#(s(x))$
 b. $\text{penult}(x) \stackrel{\text{def}}{=} \text{final}(s(x))$

The formula in (34a) defines $\text{final}(x)$ to be true just in case its successor is the word boundary #. From this, we can define the penult, which is simply an x whose successor is final. Thus, (34b) defines $\text{penult}(x)$ as $\text{final}(s(x))$, which is true if and only if the successor of x ($s(x)$) is final.

To define the unbounded shift function in Zigula, we can define the output association relation using these two predicates as follows.

- (35) $A'(x, y) \stackrel{\text{def}}{=} H(x) \wedge \text{final}(x) \wedge \sigma(y) \wedge \text{penult}(y)$

That is, $A'(x, y)$ is true if and only if x is the final H tone and y is the penultimate syllable. This shifts the association of the final H tone to the penult, regardless of where it was associated in the input. An example is given in (36). In this example, $A'(x, y)$ is only true for the H tone syllable 5, so only these two elements are associated in the output.

- (36)
$$\begin{array}{ccc} \# \text{ H} & & \# \text{ H} & & \# \\ | & & \text{---} & & \\ \# \sigma \sigma \sigma \sigma \sigma \# & \mapsto & \# \sigma \sigma \sigma \sigma \sigma \# \\ 1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7 & & 1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7 \end{array}$$

As an analysis that directly relates the input to the output, this is reminiscent of the OT conception of the process (as in, e.g., Bickmore 1996; Yip 2002), in which an ALIGN constraint pulls the H tone to the right. However, the A-ISL analysis of Zigula is local over ARs *without* reference to metrical structure per se, although it does directly reference a particular position. (Note that we could define it metrically with a predicate $\text{accented}(y)$ by using this in place of $\text{penult}(y)$ in (35), and thus this also would be A-ISL.)

More importantly, this is our first case of a non-local pattern—formally, not an ISL string function—that is local over ARs. In other words this is the first case we have reviewed in which ARs render a non-local pattern local in the computational sense advocated for here. The reason for this is that the definition of the output association relation $A'(x, y)$ only depends on information that is local to x and y . That is, as shown in (34), identifying whether a position is final or penultimate only requires information that is local to that position.

Interestingly, this is not the case for unbounded spread, in which an underlying H tone spreads in a particular direction until it is blocked or it reaches a certain position in the word. In Shambaa (Odden, 1982), an H tone spreads to the right until it reaches the penult.


- (37) Shambaa (Odden, 1982)
- | | | | |
|----|-------------------------------|-------------------------------|----------------------------------|
| a. | /ku-hand-a/ | [ku-hand-a] | 'to plant' |
| b. | /ku-fúmbatíf-a/ | [ku-fúmbátíf-a] | 'to tie securely' |
| c. | /ku-hand-ij-an-a/ | [ku-hand-ij-an-a] | 'to plant for each other' |
| d. | /ku-fúmbatíf-ij-an-a/ | [ku-fúmbátíf-íj-án-a] | 'to tie securely for each other' |
| e. | /ku-funt ^h -a/ | [ku-funt ^h -a] | 'to wash' |
| f. | /ku-tfí-funt ^h -a/ | [ku-tfí-fúnt ^h -a] | 'to wash' |
| g. | /ku-ɣofo-a-ɣofo-a/ | [ku-ɣofo-a-ɣofo-a] | 'to do repeatedly' |
| h. | /ku-tfí-ɣofo-a-ɣofo-a/ | [ku-tfí-ɣófo-á-ɣófo-a] | 'to do repeatedly' |

This is not ISL for the same reason as unbounded shift. Consider the penultimate syllables in each of the following forms.

- (38) a. $\sigma\sigma\sigma\sigma \mapsto \sigma\sigma\sigma\sigma$
 b. $\acute{\sigma}\acute{\sigma}\acute{\sigma}\acute{\sigma} \mapsto \acute{\sigma}\acute{\sigma}\acute{\sigma}\acute{\sigma}$
 $\sigma\acute{\sigma}\sigma\acute{\sigma} \mapsto \sigma\acute{\sigma}\acute{\sigma}\acute{\sigma}$


Just as with unbounded shift in (31), whether the penult surfaces as σ or $\acute{\sigma}$ depends on a trigger that may be any distance to the right. Thus, there is no QF definition for this process, and so it is not ISL. This aligns with the intuition that it is a long-distance process.

In contrast to unbounded shift, however, unbounded spread is also not A-ISL, and thus ARs do not make this pattern local. The reason is that it is not just the penult that is affected by a triggering H-tone, but *all* TBUs intervening between the trigger and the penult. As an autosegmental process, we already saw unbounded spread as one part of a possible analysis as above in (33).




The function described by the iterative rule in (39a), as exemplified in (39b), associates a single H to all following TBUs, with the exception of the penult. Defining this process in logical terms again involves defining the output association relation, $A'(x, y)$. Note, however, that this involves identifying as y *all* TBUs following the TBU associated to the H in the input. This is not QF-definable.

To see why, consider the formula for the bounded spread originally given in (25), repeated below in (40a). As shown in the example in (40b), this spreads the H tone one TBU to the right, but does not go any further.



In (40b) the H and syllable 2 satisfy $A(x, p(y))$ and so are associated in the output. However, the H and syllable 3 do not satisfy either disjunct, and so they are not associated. In this case then spread cannot continue all the way to the penult.

We can attempt to fix this problem by adding a third disjunct $A(x, p(p(y)))$, which is true for an x and the second TBU following its underlyingly associated TBU. This formula is given in (41a).



As shown in the example in (41b), this new definition associates an H with the second TBU to its right (syllable 3 in this case). However, it incorrectly would miss a *third* TBU to its right (syllable 4 in this case). In general, we can keep adding disjuncts to our QF formula, but there is always some form of enough length such that this formula will not create associations all the way to the penult.⁷

Thus, unbounded spread is not A-ISL. This is important for two reasons. One, it is an example of a non-local pattern that is *not* made local by ARs. Two, this establishes a distinction between this pattern and unbounded shift. As seen in (33), some analyses have described unbounded shift as the composition of unbounded spread and delinking. Under such an analysis, unbounded spread is the more basic operation. However, a rigorous definition of locality shows exactly the opposite: unbounded shift is local, because it only refers to information that is within some fixed distance to each element being associated, whereas unbounded spread is not, because it requires information that can be unboundedly far away from a potentially associated TBU.

⁷We cannot associate the H to *all* non-final TBUs, because this will incorrectly spread to the left as well.

3.2 Variants of Meussen’s rule

Next we consider variants of Meussen’s rule, in which an H tone is deleted or changed to L following another H tone. Surprisingly, not all of these are A-ISL. First, in Arusa, the last H tone in a phrase is deleted following another H tone, no matter the distance (Levergood, 1987; Odden, 1994).

- (42) Arusa (Levergood, 1987, p. 58)
- | | | | |
|----|-------------------|-------------------|-----------------------|
| a. | /sídáy/ | [sídáy] | ‘good’ |
| b. | /enkér sídáy/ | [enkér sídáy] | ‘good chair’ |
| c. | /olórika sídáy/ | [olórika sídáy] | ‘good ewe’ |
| d. | /kití/ | [kití] | ‘small’ |
| e. | /ádól enkér kití/ | [ádól enkér kití] | ‘I see the small ewe’ |

In (42a), the underlying high toned TBUs in /sídáy/ ‘good’ surface faithfully, whereas in (42b) and (42c) they are deleted. This can be attributed to the existence of H-tones in the preceding words. Note, however, that this deletion applies at a distance, as the triggering H-tone in (42c) is two syllables away from the target. (That the H tones disappear on both TBUs can be explained autosegmentally by positing that they are associated to the same H tone, as will be discussed in a moment.)

That this process is long-distance can be confirmed formally by observing that it is not ISL. Restated as a map over TBUs, unbounded H deletion in Arusa is exemplified schematically in (43).

- (43)
- | | |
|----|---|
| a. | $\sigma\sigma\sigma\acute{\sigma} \mapsto \sigma\sigma\sigma\sigma$ |
| b. | $\sigma\acute{\sigma}\sigma\acute{\sigma} \mapsto \sigma\acute{\sigma}\sigma\sigma$ |
| c. | $\acute{\sigma}\sigma\sigma\acute{\sigma} \mapsto \acute{\sigma}\sigma\sigma\sigma$ |

To see why this map is not ISL, consider the definition for a predicate $\acute{\sigma}'(x)$ that specifies exactly when a TBU surfaces as H-toned in the output. Compare the final syllable in (43a), which surfaces as $\acute{\sigma}$, with those in (43b) and (43c), which surface as σ because they are preceded by $\acute{\sigma}$ TBUs. In (43b), this $\acute{\sigma}$ TBU is one syllable away, whereas in (43c) it is two syllables away. To capture exactly these two examples, we could define $\acute{\sigma}'(x)$ as $\acute{\sigma}(x) \wedge \neg(\acute{\sigma}(p(p(x))) \vee \acute{\sigma}(p(p(p(x))))))$, which states that x is H-toned in the output when it is H-toned in the input *and* neither the TBU two TBUs to the left is H-toned ($\acute{\sigma}(p(p(x)))$) nor is the TBU three TBUs to the left ($\acute{\sigma}(p(p(p(x))))$). However, just as in unbounded shifting, this will be inadequate in the general case: given that the trigger can be any number of TBUs to the left, there is no finite number of $\acute{\sigma}(p(\dots p(x)\dots))$ statements that would capture all of the relevant cases in which a deletion trigger appeared. Thus, the Arusa map is not QF-definable over strings and therefore not ISL.

However, it is A-ISL. Autosegmentally, the H-deletion process in Arusa can be described with the rule in (44a).

- (44)
- | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|----|---|---|---|-----------|---|-----------|---|--|--|--|---|--|--|--|---|---|---|--|---|--|---|---|---|--|---|--|---|---|---|--|---|
| a. | $H \rightarrow \emptyset / H \text{ ____ } \#$ | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| b. | <table border="0" style="margin-left: 40px;"> <tr> <td></td> <td style="text-align: center;">H</td> <td></td> <td style="text-align: center;">H</td> <td style="text-align: center;">\mapsto</td> <td style="text-align: center;">H</td> </tr> <tr> <td></td> <td style="text-align: center;"> </td> <td></td> <td style="text-align: center;">/</td> <td></td> <td style="text-align: center;"> </td> </tr> <tr> <td></td> <td style="text-align: center;">σ</td> <td style="text-align: center;">σ</td> <td style="text-align: center;">σ</td> <td></td> <td style="text-align: center;">σ</td> </tr> <tr> <td></td> <td style="text-align: center;">σ</td> <td style="text-align: center;">σ</td> <td style="text-align: center;">σ</td> <td></td> <td style="text-align: center;">σ</td> </tr> <tr> <td></td> <td style="text-align: center;">σ</td> <td style="text-align: center;">σ</td> <td style="text-align: center;">σ</td> <td></td> <td style="text-align: center;">σ</td> </tr> </table> | | H | | H | \mapsto | H | | | | / | | | | σ | σ | σ | | σ | | σ | σ | σ | | σ | | σ | σ | σ | | σ |
| | H | | H | \mapsto | H | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | | / | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | σ | σ | σ | | σ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | σ | σ | σ | | σ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | σ | σ | σ | | σ | | | | | | | | | | | | | | | | | | | | | | | | | | |

The rule in (44a) operates as shown in (44b): a final H tone deletes immediately following another H. Note that this operates over the *tonal* tier, and so the rule applies even to Hs associated to non-adjacent TBUs.

This locality on the tier means that this process is A-ISL. As described in §2, logical definitions specify deletion via an input element that does not satisfy any formula for the output labels. To capture deletion in Arusa, then, we write a formula for output H nodes such that it is not satisfied by any input H that is both final and is preceded by another H. This is given in (45).

$$(45) \quad H'(x) \stackrel{\text{def}}{=} H(x) \wedge (\text{final}(x) \rightarrow \neg H(p(x)))$$

This formula is interpreted as follows. The output label of x is H if and only if it satisfies two conditions. First, x must be an H in the input ($H(x)$). Second, if x is the final member of its tier, it cannot be preceded by another H ($\text{final}(x) \rightarrow \neg H(p(x))$). For example, the first H in (44b) satisfies both conditions—it is an H, and it is not final (so it satisfies the implication by failing its antecedent). The second, however, does not satisfy both conditions—it is H, but it is last and it is preceded by another H. Thus, it does not satisfy $H'(x)$ and will not receive a label in the output and so is deleted by the convention that unlabeled output elements do not surface.

Thus, long-distance deletion in Arusa is A-ISL, while it is not ISL. Like unbounded shift (as discussed in §3.1), it is an example of a non-local process that, by our computational metric, is made local by ARs. However, counterintuitively, bounded variants of Meussen’s Rule are not A-ISL. There are two ways in which this can happen. One, modification of the H tone crucially refers to information on another tier, which is non-local according to the definition of A-ISL. Two, modification of the H depends on the *output* of the preceding tone, which is not input-oriented as required by A-ISL. Both variants are attested.

First, we consider the case of Luganda, in which timing tier information is crucial in determining whether Meussen’s rule applies. In Luganda (Hyman and Katamba, 2010), an underlying H tone lowers to L (marked below with a grave accent) immediately following another underlying H tone. An example is shown in (46). The output forms given below are intermediate, before the application of other processes; see Hyman and Katamba (2010).

- (46) Luganda (Hyman and Katamba, 2010)
- | | | | |
|----|-------------------------|-----------------------|--------------------------------|
| a. | /a-láb-a/ | a-láb-a | ‘s/he sees’ |
| b. | /bá-láb-a/ | bá-làb-a | ‘they see’ |
| c. | /bá-lí-láb-a/ | bá-lì-làb-a | ‘they will see’ |
| d. | /a-bá-tá-lí-láb-il-ila/ | a-bá-tà-lì-làb-il-ila | ‘they who will not look after’ |
| e. | /bá-ki-láb-a/ | bá-ki-láb-a | ‘they see it’ |

The contrast between (46b) and (46e) illustrate the locality condition of Meussen’s Rule in Luganda. In (46b), the /á/ in the root /-láb-/ immediately follows the prefix H in /bá-/ and so is lowered. In contrast, in (46e) the root H is separated from the prefix H by the intervening toneless TBU /ki/, and thus is not lowered.

As a string map, this process is ISL. Examples of Meussen’s Rule in Luganda as changes in strings of TBUs are given in (47).

- (47)
- | | |
|----|---|
| a. | $\sigma\acute{\sigma}\sigma \mapsto \sigma\acute{\sigma}\sigma$ |
| b. | $\acute{\sigma}\acute{\sigma}\sigma \mapsto \acute{\sigma}\grave{\sigma}\sigma$ |
| c. | $\acute{\sigma}\acute{\sigma}\acute{\sigma}\sigma \mapsto \acute{\sigma}\grave{\sigma}\grave{\sigma}\sigma$ |
| d. | $\acute{\sigma}\acute{\sigma}\acute{\sigma}\sigma \mapsto \sigma\acute{\sigma}\sigma\acute{\sigma}$ |

Capturing this logically requires defining two predicates, $\acute{\sigma}'(x)$ and $\grave{\sigma}'(x)$, that specify when a syllable is realized as an H-toned or L-toned, respectively.

- (48)
- | | |
|----|---|
| a. | $\acute{\sigma}'(x) \stackrel{\text{def}}{=} \acute{\sigma}(x) \wedge \neg\acute{\sigma}(p(x))$ |
| b. | $\grave{\sigma}'(x) \stackrel{\text{def}}{=} \acute{\sigma}(x) \wedge \acute{\sigma}(p(x))$ |

These formulas are interpreted as follows. In (48a), $\acute{\sigma}'(x)$ specifies that a TBU is output as H-toned if and only if it is H-toned in the input ($\acute{\sigma}(x)$) and its predecessor is not H-toned in the input ($\neg\acute{\sigma}(p(x))$). This captures the generalization that an H tone only surfaces as H when its predecessor is not H-toned. An example with indexed TBUs is given below. In (49) TBUs 1 and 4 satisfy $\acute{\sigma}'(x)$, as they satisfy both $\acute{\sigma}(x)$ and $\neg\acute{\sigma}(p(x))$ —1 because it has no predecessor, and 4 because its predecessor is σ . Thus, 1' and 4' are labeled $\acute{\sigma}$.

- (49)
- | | | | |
|--|---|-----------|---|
| | $\acute{\sigma} \acute{\sigma} \sigma \acute{\sigma}$ | \mapsto | $\acute{\sigma} \grave{\sigma} \sigma \acute{\sigma}$ |
| | 1 2 3 4 | | 1' 2' 3' 4' |

In contrast, TBU 2 fails $\acute{\sigma}'(x)$: it is $\acute{\sigma}$ in the input, but its predecessor (1) is also $\acute{\sigma}$, so it fails $\neg\acute{\sigma}(p(x))$. Instead, it satisfies (48b) $\grave{\sigma}'(x)$, which specifies the complementary situation: when x is H-toned in the input and its predecessor is also H ($\acute{\sigma}(p(x))$). Thus, 2' (and only 2') instead surfaces as $\grave{\sigma}$. Thus, Meussen’s Rule in Luganda is QF-definable and therefore ISL.

This is not the case over ARs. Over ARs, the desired map is as exemplified in (50).

- (50)
- | | | | |
|----|------------------------|-----------|------------------------|
| a. | H H | | H L |
| | | \mapsto | |
| | $\sigma \sigma \sigma$ | | $\sigma \sigma \sigma$ |

$$\text{b.} \quad \begin{array}{ccccccc} & \text{H} & & \text{H} & \mapsto & \text{H} & \text{H} \\ & | & & | & & | & | \\ \sigma & \sigma & \sigma & & & \sigma & \sigma & \sigma \end{array}$$

The crucial difference is between the second H in (50a) and the second H in (50b): both follow an H in the input, but only in (50a) is the preceding H associated to an immediately preceding TBU. Thus, the second H in (50a) lowers to L, while that in (50b) does not.

We can attempt to capture this in a similar way as we did over strings, by defining $H'(x)$ such that it is only true when H does not follow another H, and the converse for $L'(x)$:

$$(51) \quad \begin{array}{l} \text{a.} \quad H'(x) \stackrel{\text{def}}{=} H(x) \wedge \neg H(p(x)) \\ \text{b.} \quad L'(x) \stackrel{\text{def}}{=} H(x) \wedge H(p(x)) \end{array}$$

However, this does not achieve the desired effect, as (52) shows.

$$(52) \quad \text{a.} \quad \begin{array}{ccccccc} & \text{H} & \text{H} & & \mapsto & \text{H} & \text{L} \\ & | & | & & & | & | \\ \sigma & \sigma & \sigma & & & \sigma & \sigma & \sigma \end{array}$$

$$\text{b.} \quad \begin{array}{ccccccc} & \text{H} & & \text{H} & \mapsto & * & \text{H} & \text{L} \\ & | & & | & & & | & | \\ \sigma & \sigma & \sigma & & & & \sigma & \sigma & \sigma \end{array}$$

The problem is that the second H tone in both (52a) and (52b) follows another H tone in the input. Thus, they both are output as L, which is incorrect in the case of (52b). The reason this doesn't work is the same reason that the long-distance version of Meussen's Rule in Arusa discussed above *does* work: whether one H tone is the predecessor of another is completely unrelated to the number of TBUs between them.

Here we recall the theorem in (19), repeated in (53), which asserts that for a process to be A-ISL, any changes on a particular tier must *only* use information that is local in the input *on that tier*.

$$(53) \quad \textbf{Theorem.} \text{ If an AR map is A-ISL, then the individual map on each tier is ISL. (Chandlee and Jardine, 2019a, Thm 1)}$$

The Luganda map is not local in this way; whether or not an H changes to L depends both on whether or not there is a preceding H *and* whether their associated TBUs are adjacent on the TBU tier. However, checking this requires quantification, as illustrated below in (54).

$$(54) \quad L'(x) \stackrel{\text{def}}{=} H(x) \wedge H(p(x)) \wedge (\exists y)[A(x, y) \wedge A(p(x), p(y))]$$

The definition in (54) is identical to that in (51b) except that it adds a third condition, namely that there is some TBU y such that x is associated to y and also that $p(x)$ is associated to $p(y)$. This ensures that x and its preceding H $p(x)$ are associated to adjacent TBUs. However, this statement uses an existential quantifier \exists , and so is not QF.

Meussen's Rule in Luganda fails to show the property in (53), because the map on the tone tier is not even a function: a tone tier of, e.g., HH can be mapped to either HL (when the two Hs are associated to adjacent TBUs) or HH (when they are associated to nonadjacent TBUs). Thus, it is not A-ISL. This is significant because it shows there are maps that are ISL but not A-ISL, or in other words, there are local maps that ARs actually render *non-local*.

We can also use (53) to show that another seemingly local variant of Meussen's Rule is neither ISL nor A-ISL. As in Luganda, in Shona, Meussen's Rule lowers an H to an L following another H (Odden, 1986; Myers, 1987). However, in Shona, as opposed to Luganda, an H tone to which Meussen's Rule has applied does not serve as a trigger for a following H tone. This results in an alternating pattern of H- and L-toned TBUs.

$$(55) \quad \text{Shona (Odden, 1986)}$$

- a. /hóvé/ [hóvé] ‘fish’
 b. /né-hóvé/ [né-hòvè] ‘with-fish’
 c. /né-é-hóvé/ [né-è-hóvé] ‘with-of-fish’
 d. /né-é-é-hóvé/ [né-è-é-hòvè] ‘like-with-of-fish’

In (55), the underlying H tones in /hóvé/ ‘fish’ surface as L when following the H tone in the prefix /né-/ ‘with’ in (55b). That both TBUs lower can be explained by positing that they are associated to the same H tone; this is discussed in more detail below. More importantly, this lowering is blocked when the H-toned prefix /é-/ ‘of’ intervenes in (55c). Instead, the H tone of the prefix lowers.

This shows that the correct generalization for Shona is that Meussen’s Rule lowers an H to an L following another H tone *in the output*. Thus, because the second H in (55c) itself lowers by Meussen’s Rule, it cannot become the trigger for the H TBUs in /hóvé/. That this creates an alternating pattern of H- and L-toned TBUs is illustrated dramatically in (55d), in which four successive H tones surface as HLHL.

Because an H tone is lowered to L only following another *output* H, this map is not ISL. Simplifying somewhat, the map for Meussen’s Rule in Shona is as follows, in which a string of H-toned TBUs surfaces as a string of alternating H- and L-toned TBUs.

- (56) a. $\acute{\sigma}\acute{\sigma} \mapsto \acute{\sigma}\grave{\sigma}$
 b. $\acute{\sigma}\acute{\sigma}\acute{\sigma} \mapsto \acute{\sigma}\grave{\sigma}\acute{\sigma}$
 c. $\acute{\sigma}\acute{\sigma}\acute{\sigma}\acute{\sigma} \mapsto \acute{\sigma}\grave{\sigma}\acute{\sigma}\grave{\sigma}$

To see why this map is not ISL, consider what properties of an input TBU would need to be identified for $\acute{\sigma}'(x)$ and $\grave{\sigma}'(x)$. The issue is best illustrated diagrammatically. The following contrasts the outputs for the second, third, and fourth H-toned syllables in example (56c).

- (57) a. $/ \acute{\sigma} \acute{\sigma} \boxed{\acute{\sigma}} \acute{\sigma} /$
 $[\acute{\sigma} \grave{\sigma} \boxed{\acute{\sigma}} \grave{\sigma}]$
 b. $/ \acute{\sigma} \boxed{\acute{\sigma}} \acute{\sigma} \acute{\sigma} /$
 $[\acute{\sigma} \boxed{\grave{\sigma}} \acute{\sigma} \grave{\sigma}]$

What conditions do the third output syllable satisfy that lead to it bearing an H tone? In terms of the input, it has an H tone and its predecessor and successor are also H-toned ($\acute{\sigma}(x) \wedge \acute{\sigma}(p(x)) \wedge \acute{\sigma}(s(x))$). But the exact same conditions hold for the second output syllable (shown in (57b)), which instead is L-toned. We can try expanding the window of input TBUs to the left or right, but this will be of no use—clearly, with an input string of H-toned syllables no amount of information will be enough to distinguish these two output positions. What is needed is reference to the output, particularly whether or not the previous output syllable bears H or L.

For the same reason, this alternating Meussen’s Rule pattern in Shona is not A-ISL. As a map over ARs, the alternating application of lowering occurs over H tones on the tonal tier. This captures why both TBUs in /hóvé/ ‘fish’ lower in (55b) [né-hòvè] ‘with-fish’ and (55d) [né-è-é-hòvè] ‘like-with-of-fish’, as we can represent these two TBUs as associated to the same underlying H tone. This is shown below in (58), which represents (55b) and (55d) as AR maps.

- (58) a. $\begin{array}{ccc} H & H & \\ | & \diagdown & \\ \sigma & \sigma & \sigma \end{array} \mapsto \begin{array}{cc} H & L \\ | & \diagdown \\ \sigma & \sigma \end{array}$
 b. $\begin{array}{cccc} H & H & H & H \\ | & | & | & \diagdown \\ \sigma & \sigma & \sigma & \sigma \end{array} \mapsto \begin{array}{cccc} H & L & H & L \\ | & | & | & \diagdown \\ \sigma & \sigma & \sigma & \sigma \end{array}$

Recasting the process as a map over ARs, however, does not change the fundamental property that made it not QF-definable over strings. The AR map must map a string of Hs to an alternating string of Hs and Ls. For example, (59) contrasts two input H tones in (58b), focusing on the tonal tier.

- (59) a. / HH[H]H /
 [HL[H]L]
 b. / [H]HH /
 [H[L]HL]

Just as in (57), there is no local input information that distinguishes the third H in the tonal tier, which does not lower to L (59a), and the second H, which does (59b). Thus, the map on the tonal tier is not ISL. Therefore, as per the theorem in (19), the map is also not A-ISL.

3.3 Interim summary

We have studied a number of tonal processes with respect to whether they are ISL or A-ISL. Table 1 summarizes the results.

Pattern	ISL	A-ISL
Bounded shift	✓	✓
Bounded spread	✓	✓
Unbounded shift	✗	✓
Unbounded spread	✗	✗
Unbounded Meussen’s Rule	✗	✓
Bounded Meussen’s Rule	✓	✗
Alternating Meussen’s Rule	✗	✗

Table 1: Summary of analyses

Several points are of interest. Perhaps unsurprisingly, bounded shift and spread are both ISL and A-ISL. Unbounded shift, while not ISL, was A-ISL, which gives an example of a non-local process that is local when represented with ARs. Again, this was due to the fact that the asynchronous nature of ARs allow distant elements to be connected through local manipulation of the association relation. In contrast, unbounded spread is neither, because changes are based on information that is not local in the input. It is, however, based on information that is local in the output—this is discussed more in §5.

In terms of modification of tones on the tonal tier, unbounded Meussen’s Rule in Arusa is A-ISL, whereas it is not ISL, providing another example of a non-local process that is made local by ARs. This is because the triggering H tone is local to the target H tone on the tonal tier, even though the two tones may be associated to distant TBUs. In contrast, bounded Meussen’s Rule in Luganda is not A-ISL, because both information on the tonal tier and on the TBU tier is necessary to determine whether or not the rule applies. Finally, the ‘alternating’ version of Meussen’s found in Shona is also not A-ISL. As discussed below in §5, this process, like unbounded spread, is output-local.

Thus far we have been analyzing processes in isolation, as a means of investigating locality and ARs on a case-by-case basis. But of course processes often interact with other processes or factors, and it is equally worthwhile to investigate those cases to see if the interaction affects classification as ISL/A-ISL. To that end, in the next section we turn to a more detailed analysis of multiple processes in Shona.

4 A detailed analysis: OCP effects in Shona

We have already concluded that Meussen’s rule in Shona is not A-ISL. However, two related processes analyzed by Myers (1987, 1997) as being motivated by the OCP are. Specifically, *tone spread* and *tone slip* find a unified analysis under a logical definition of association.

The OCP condition described in Myers (1997) forbids the following configuration, in which two adjacent H tones are associated to adjacent syllables.

$$(60) \quad \begin{array}{cc} \text{H} & \text{H} \\ | & | \\ \sigma & \sigma \end{array}$$

In logical terms, we can describe this configuration as follows. First we define a formula $\text{next}_A(x, y)$ which is true for x and y when the successor of x is associated to the successor of y .

$$(61) \quad \text{next}_A(x, y) \stackrel{\text{def}}{=} A(s(x), s(y))$$

Two elements x and y are in an OCP configuration if they satisfy the following formula.

$$(62) \quad A(x, y) \wedge \text{next}_A(x, y)$$

This identifies an x and y such that x is associated to y , and furthermore that the distinct successor of x is associated to the distinct successor of y . We will now see how this factors in the definition of association in the phonological word in Shona.

First, in tone spread, an underlying tone spreads to the subsequent unspecified TBU across a morphological boundary within a phonological word, as shown in (63).

$$(63) \quad \begin{array}{lll} \text{a.} & /sadz\acute{a}/ & [sadz\acute{a}] \quad \text{'porridge'} \\ & /i\text{-sadz\acute{a}}/ & [i\text{-s\acute{a}dz\acute{a}}] \quad \text{'(it) is porridge'} \\ \text{b.} & /ku\text{-verenga}/ & [ku\text{-verenga}] \quad \text{'read (inf.)'} \\ & /ti\text{-ch\acute{a}\text{-verenga}}/ & [ti\text{-ch\acute{a}\text{-v\acute{e}renga}] \quad \text{'we will read'} \end{array}$$

However, the spreading is blocked when there is an H tone associated to the TBU following the target TBU.

$$(64) \quad \begin{array}{lll} /badz\acute{a}/ & badz\acute{a} & \text{'hoe'} \\ /i\text{-badz\acute{a}}/ & i\text{-badz\acute{a}} & \text{'(it) is a hoe'} \\ *i\text{-b\acute{a}dz\acute{a}} & & \end{array}$$

This can be attributed to tone spread being blocked just in case it creates the configuration in (60). In terms of ARs, the transformations in (63a) and (64) are as below in (65a) and (65b), respectively.

$$(65) \quad \begin{array}{ll} \text{a.} & \begin{array}{ccc} \text{H} & & \text{H} \\ | & & \diagdown \\ \sigma & \sigma & \sigma \end{array} \mapsto \begin{array}{ccc} & & \text{H} \\ & & \diagdown \\ \sigma & \sigma & \sigma \end{array} \\ \\ \text{b.} & \begin{array}{ccc} \text{H} & \text{H} & \\ | & | & \\ \sigma & \sigma & \sigma \end{array} \mapsto \begin{array}{ccc} \text{H} & \text{H} & \\ | & | & \\ \sigma & \sigma & \sigma \end{array}, \quad * \begin{array}{ccc} \text{H} & \text{H} & \\ \diagdown & | & \\ \sigma & \sigma & \sigma \end{array} \end{array}$$

In a related process of tone *slip*, a tone underlyingly associated to multiple TBUs will retract one TBU when immediately preceding another H-tone span.

$$(66) \quad \begin{array}{lll} /b\acute{a}ng\acute{a}/ & [b\acute{a}ng\acute{a}] & \text{'knife'} \\ /g\acute{u}r\acute{u}/ & [g\acute{u}r\acute{u}] & \text{'big'} \\ /b\acute{a}ng\acute{a} g\acute{u}r\acute{u}/ & [b\acute{a}ng\acute{a} g\acute{u}r\acute{u}] & \text{'big knife'} \end{array}$$

Again, this is to avoid the structure specified in (60). Autosegmentally, the transformation of underlying /b\acute{a}ng\acute{a} g\acute{u}r\acute{u}/ 'big knife' to [b\acute{a}ng\acute{a} g\acute{u}r\acute{u}] is as follows.

$$(67) \quad \begin{array}{ccc} \text{H} & \text{H} & \\ \diagdown & \diagdown & \\ \sigma & \sigma & \sigma \end{array} \mapsto \begin{array}{ccc} \text{H} & \text{H} & \\ | & \diagdown & \\ \sigma & \sigma & \sigma \end{array}$$

In sum, an underlying H tone spreads unless it creates an OCP violation. Similarly, if there is an underlying association that violates the OCP, the first H tone retracts to avoid the violation on the surface.

A logical definition of output association treats these two processes in a unified way. First, in neither case does an association surface when it creates an OCP violation. Any x and y associated in the output cannot be in an OCP configuration; thus they must satisfy the following, based on the logical characterization of an OCP violation given in (62).


$$(68) \quad \neg(\text{next}_A(x, y))$$

The above formula is only satisfied for x and y whose successors are *not* associated. This is essentially the negation of the structure in (62), except that (68) is agnostic to whether x and y are associated in the input.

The following is then true about a TBU y in the output: it is associated to a tone x if x was *associated to y 's predecessor in the input*, and if x and y adhere to the above OCP condition in (68). This is stated in the formula below in (69).

$$(69) \quad A(x, p(y)) \wedge \neg(\text{next}_A(x, y))$$

This formula is true for exactly the set of output tone-TBU pairs highlighted with dashed lines in the below forms, taken from (65) and (67) above.



Note that this formula satisfies a pair of TBUs in the spreading context, as in (70a), as well as a pair of underlying TBUs that are not in an OCP-violating context, as in (70c). Conversely, pairs in the OCP-violating context, such as the first H and second TBU in both (70b) and (70c), will not satisfy the formula in (69), regardless of whether they were spreading targets (as in (70b)) or were underlyingly associated (as in (70c)). Thus, the logical statement in (69) captures both tone slip and the blocking of tone spread.

To define the full output association relation, however, we must add one final clause, as the above does not hold for an H tone and the first TBU to which it is associated in the input. First, we define the following predicate, which is true when y is the first TBU associated to a tone x in the input.

$$(71) \quad \text{first}_A(x, y) \stackrel{\text{def}}{=} A(x, y) \wedge \neg A(x, p(y))$$

The statement $A(x, p(y))$ is true when there is a distinct predecessor of y that is associated to x ; in other words, there is some previous TBU that is associated to x . The full predicate $\text{first}_A(x, y)$ is thus true only when $A(x, y)$ and there is no such $p(y)$.


The full output association relation for Shona can thus be defined as the following disjunction.

$$(72) \quad A'(x, y) \stackrel{\text{def}}{=} \text{first}_A(x, y) \vee (A(x, p(y)) \wedge \neg \text{next}_A(x, y))$$

The first disjunct in (72) copies over an input association from x to y just in case y is the first TBU associated to x in the input. The second disjunct is exactly the formula from (69), and thus implements tone spread and tone slip such that surface violations of the OCP are avoided. That is, it is true for any x and y such that x is associated to y 's predecessor in the input, so long as x and y are not in a potential OCP-violating configuration.

This completes the map except for one case: when an H tone is associated to a single TBU in an OCP

violation. In this case, the definition in (72) maps this configuration faithfully.



This is in fact the exact case in which Meussen’s rule applies, as seen in the following examples.

- (74)
- | | | | |
|----|-------------|-------------|----------------------|
| a. | /bángá/ | [bángá] | ‘knife’ |
| | /í-bángá/ | [í-banga] | ‘(it) is a knife’ |
| b. | /sékúru/ | [sékúru] | ‘grandfather’ |
| | /vá-sékúru/ | [vá-sekuru] | ‘grandfather (hon.)’ |

As established in §3, this version of Meussen’s Rule is not A-ISL. However, we can capture the application of this process by composing the A-ISL tone spread/slip function with the one for Meussen’s Rule (whatever its ultimate characterization). Thus, Meussen’s Rule will apply just in case the tone spread/slip definition leaves the configuration in (73).

5 Discussion

This concludes our exploration of how the A-ISL functions, which implement a computational notion of input locality based, fare in capturing long-distance tone processes. Some processes, such as unbounded shift, are local only under this formulation, whereas other processes, such as unbounded spreading, are not. Other intuitively ‘local’ processes, such as locally bounded Meussen’s Rule, are also not A-ISL, because a change on one tier is conditioned by information on another tier.

In §5.1, we compare and contrast A-ISL with previous notions of locality invoked in phonological theorizing. We find that, while A-ISL differs in important ways from these previous notions, it does capture an important intuition that representations should make processes simpler, or more local. An important question of output-based locality, raised by Optimality Theory, is discussed in detail in §5.2.

5.1 Previous considerations of locality and representation


Considerations of locality have loomed large in the literature on phonological representations. Odden (1995, pg. 474) states:


A widely held desideratum in phonological theory—indeed much of the motivation for nonlinear phonology and one of the outstanding problems of linear phonology—is that rules should be “local.” Though there are many unresolved problems in the locality issue, it is generally agreed that a local rule formulation would only allow specification of one element to the right and/or left of the focus.

As an example, Odden raises the process of unbounded spread to antepenultimate position in the Nguni languages, such as Ndebele (Sibanda, 2004; Hyman, 2011):


- (75) Ndebele (Sibanda, 2004; Hyman, 2011)
- | | | | |
|----|---------------------|---------------------|-------------------------|
| a. | /ú-ku-hlek-a/ | [ú-kú-hlek-a] | ‘to laugh’ |
| b. | /ú-ku-hlek-is-a/ | [ú-kú-hlék-is-a] | ‘to amuse (make laugh)’ |
| c. | /ú-ku-hlek-is-an-a/ | [ú-kú-hlék-ís-an-a] | ‘to amuse each other’ |

Odden notes that it is possible to posit the iterative rule in (76) for this process. The target for the rule is a syllable that is followed by two other syllables. As shown in (77), this will iterate up until the penult, which is followed by only one syllable.





However, this rule is non-local according to Odden’s definition, because it refers to more than one syllable to the target’s right. Instead, Odden argues that an accentual analysis makes such rules local. If the antepenultimate syllable receives an accent, then we can formulate a tone-accent attraction rule, as in (78a).



This rule then draws an association between the H and the antepenult. While Odden does not explicitly state how spreading would be accomplished, one could imagine that the Well-Formedness Condition (WFC; Goldsmith, 1976) would apply and fill in the intervening TBUs, as in the derivation in (78b). The rule in (78a) is local, according to Odden (1995)’s definition, as it refers only to the target. However, there is a bit of a trick here: it creates a gapped structure in between two non-local TBUs—the initial and antepenult syllables—over which the WFC then operates.

This concern for locality has been no less important for Optimality Theory (OT; Prince and Smolensky, 2004). As McCarthy (2010) states, one of the “fundamental descriptive and explanatory goals of OT” is to “to derive complex patterns from the interaction of simple constraints” (p. 200). One such concrete proposal is the articulatory notion of ‘strict locality,’ meaning articulatory locality at the level of the segment (Gafos, 1996; Chiośain and Padgett, 2001). Under such a theory, spreading is motivated by output markedness constraints grounded by the phonetics governing the articulation of adjacent segments.

However, not all tone patterns can be motivated entirely by local output conditions. A case in point is that of bounded shift in Rimi (Myers, 1997), discussed in §§2 and 3, repeated below in (79).

- (79) Rimi (Schadeberg, 1979; Myers, 1997)
- | | | | |
|----|-------------|-------------|---------------|
| a. | /u-hang-a/ | [u-hang-a] | ‘to meet’ |
| b. | /u-púm-a/ | [u-púm-á] | ‘to go away’ |
| c. | /mu-ntu/ | [mu-ntu] | ‘person’ |
| d. | /rá-mu-ntu/ | [ra-mú-ntu] | ‘of a person’ |
| e. | /u-huvi-ĩ/ | [u-huvi-ĩ] | ‘belief’ |
| f. | /mu-tém-ĩ/ | [mu-tem-ĩ] | ‘chief’ |

Crucial to capturing this shift is the position of the underlying H tone. In order to capture this kind of behavior, Myers (1997) posits the following LOCALITY constraint.

- (80) LOCAL: (Myers, 1997, p. 876, (50)) If an input tone T has an output correspondent T' , some edge of T must correspond to some edge of T' .

Importantly, this constraint is not a markedness condition, as it crucially refers to the input. In fact, it is more of a ‘two-level’ constraint (McCarthy, 1996; Kager, 1999), in that it refers both to the input position in which a tone appears in the UR as well as the output position in which it surfaces in the SR. That this pattern is properly input local is made clear in the following section, which formalizes output locality and shows that Rimi is not output local.

Furthermore, output constraints often invoked to explain tone patterns are not ‘local’ by the definitions already put forth. Much of the foundational work on tone in OT (Myers, 1997; Cassimjee and Kisseberth, 1998; Yip, 2002) makes copious use of gradient ALIGN constraints (McCarthy and Prince, 1993, 1995).⁸ A gradient ALIGN constraint counts the distance between two elements in a representation, which is difficult

⁸Zoll (2003) shows how many tone map patterns can be explained without ALIGN constraints, but does not completely do away with them.

to fit into a conception of ‘local’. In computational terms, this kind of constraint is in fact quite complex (Eisner, 1997b). While proposals to replace ALIGN with local constraints exist (Eisner, 1997a; McCarthy, 2003), they have not been widely adopted.

However, as is clear from the analyses in §3, both input and output locality appear to be necessary to capture the full range of tonal patterns. We now discuss how this has been formalized for strings, and how it may be extended to ARs.

5.2 Logic and output locality

This paper has contrasted string versus AR analyses of tone processes using the ISL and A-ISL functions, which formalize a particular notion of locality based on the input. As is clear from the analyses in §3, this notion of locality is not sufficient for all tonal processes. It is thus worth discussing what a comprehensive theory of locality in tone might look like, and how the results of this paper would fit into it. Given the past successes of OT, it may be tempting to say that we should replace ISL and A-ISL with some notion of output locality. However, it is clear from bounded spreading and shift that any theory of tone must also include input locality. Below, we sketch out such a theory based on an extension of QF transductions that includes both input- and output-local functions. In sum, the results of this paper hold regardless of future theorizing on tone.

A well-articulated version of output locality for strings is the *output-strictly local* (OSL) functions (Chandlee, 2014; Chandlee et al., 2015b). These were originally defined in formal language and automata-theoretic terms, but we discuss them informally here. Essentially, these are the functions for which there is some k such that, for any position in the string, its output is calculated based on the previous $k - 1$ positions in the output.⁹






Figure 3: Computing an OSL function. Here, τ is an arbitrary input target; τ' is its output. The sequence $\ell_1, \ell_2, \dots, \ell_{k-1}$ in the output is the information available on the left side of τ available for computing τ' .

For example, consider unbounded spreading. Any input H-toned syllable $\acute{\sigma}$ always surfaces as $\acute{\sigma}$; whether or not an input σ does depends on whether or not the previous *output* symbol was $\acute{\sigma}$. This is illustrated below in (81).




Examples (81a) and (b) contrast two underlyingly toneless syllables (highlighted in boxes) in an underlyingly toned word versus a toneless word, respectively. In (81a), for the boxed target input σ , the preceding output position—boxed and shaded—is a $\acute{\sigma}$. The output of this target σ is then $\acute{\sigma}$. In (81b), for the boxed target input σ , its preceding output position is σ , so its own output is σ . Thus, in general, for unbounded spread the output of any position is determined by the position one syllable to the left in the output.¹⁰

In contrast, bounded shift is not OSL. The following schematically contrasts input forms in which an underlying toneless syllable directly follows a toned syllable versus a form with all toneless syllables.



⁹ $k - 1$ previous output positions are used, rather than k , because the position whose output is being determined is itself the k^{th} item in the window under consideration.


¹⁰For brevity we exclude discussion of how the final syllable is always output as σ . How this is handled does not change the fact that the basic nature of unbounded spreading is OSL.



The input syllables highlighted in boxes have differing outputs: in (82a), it is output as $\acute{\sigma}$ (the result of shift from the preceding input syllable), whereas in (82b) it is output as σ (as there is no preceding input H-toned syllable). However, as the shaded output regions in both examples show, the output information preceding each syllable is identical. As indicated by the ellipses, this information will be identical at any distance. Because there is no information within a bounded distance of the output from the targets that distinguish them, bounded shift is not OSL.

In a well-defined sense, then, bounded shift is input-local, but not output-local. This explains why Myers (1997)'s LOCALITY constraint discussed in §5.1 necessarily refers to both input and output positions: bounded spread requires that a Faithfulness violation (namely, the change in position of a tone's domain) occurs within some fixed distance from an input position. Thus, a pure notion of output locality cannot capture the full range of tonal patterns.

It is worth noting that unbounded shift is also not OSL, for the same reasons as bounded shift. This is illustrated below in (83), which contrasts a form with an underlying H tone shifting to the penult (83a) with a form with no underlying H tone (83b).



The map in (83a) shows an underlying penultimate syllable that is mapped to an H-toned syllable in the output; the map in (83b) shows an underlying penult syllable that is mapped to an unspecified syllable in the output. As in (82) for bounded shift, even though the outputs of the two targets are distinct, the preceding information to the left of them in the output is identical. Thus, unbounded shift is not OSL. Because it is A-ISL, the solution to making this process local is not to consider output locality, but to consider locality over ARs.

However, unbounded spreading shows us that we also need a notion of output-locality for ARs; that is, an extension of OSL to ARs. No such definition exists, but the work of Koser et al. (to appear) offer one way to develop such a class. They do not consider unbounded spreading per se, but instead focus on *tone map* patterns. However, their solution does extend to unbounded spreading, so we briefly discuss it here.

Tone map patterns are an important type of tonal process that is beyond QF-definability (and is thus not A-ISL). This is a pattern, such as in Mende (Leben, 1973, 1978), which can be explained by the map of an underlyingly unassociated autosegmental melody to a string of symbols. The examples below illustrate the distribution of tone in Mende.

- (84) Mende (Leben, 1973)
- | | | | | | |
|--------|-------------|----------|---------|-----------|-------------|
| a. kó | 'war' | b. pélé | 'house' | c. háwámá | 'waist' |
| d. mbú | 'owl' | e. ngílà | 'dog' | f. fèlàrà | 'junction' |
| g. mbã | 'companion' | h. nyàhà | 'woman' | i. níkìlì | 'groundnut' |

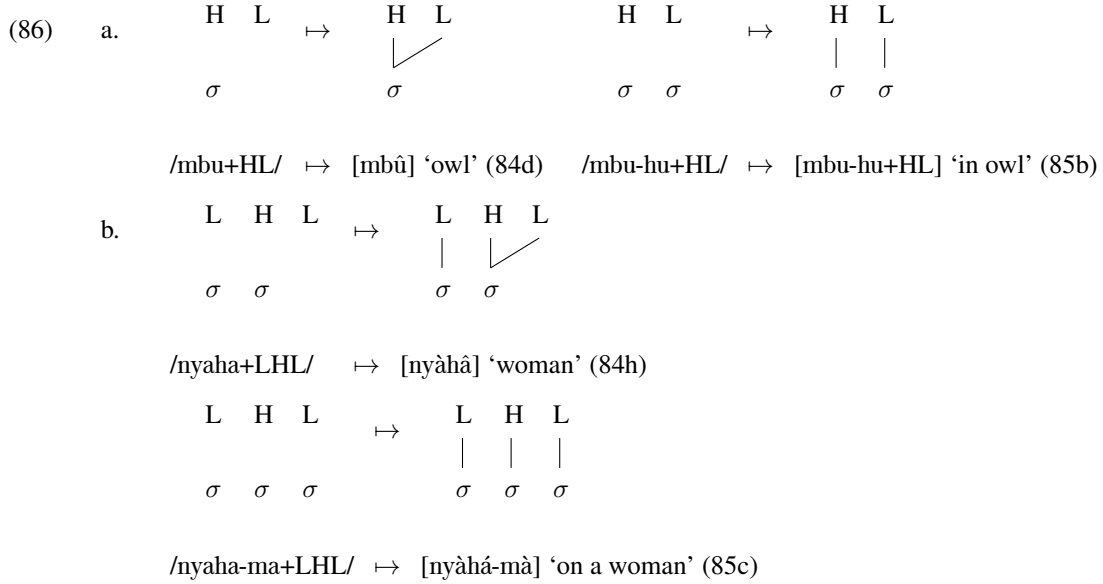
Words in Mende choose between one of several melodies, which are then distributed over the syllables in the word. For example, the words in the first row in (84) have an H melody, and thus all syllables are pronounced with a high tone (e.g., (84b) [pélé] 'house'). The words in the second row follow an HL melody; this is realized as a single falling tone in (84d) [mbú] 'owl' but as an H-L-L sequence in (84f) [fèlàrà] 'junction'. This distribution of the tone also extends to toneless suffixes, as shown below in (85).

- (85) Mende suffixes (Leben, 1973)
- | | |
|-------------|--------------|
| a. kó-hú | 'in war' |
| b. mbú-hù | 'in owl' |
| c. nyàhá-mà | 'on a woman' |

The toneless suffixes /-hu/ 'in' and /-ma/ 'on' serve as extra syllables over which the tone melody is realized: (85b) [mbú-hù] 'in owl' is an H-L sequence whereas (84d) [mbú] 'owl' is pronounced as a falling-toned syllable in isolation. Likewise, whereas the LHL-melody word (84h) [nyàhà] 'woman' is pronounced as a

low-falling sequence of syllables in isolation, (85c) [nyàhá-mà] ‘on a woman’ is pronounced as a L-H-L sequence of syllables (parallel to the three-syllable word (84i) [nikíli] ‘groundnut’).

One of the initial insights of autosegmental theory was that these patterns were the map of an underlying, unassociated melody to a span of syllables, as illustrated below in (86).




The examples in (86) show that the famous right-to-left, one-to-one map of tonal autosegments to syllables produces the right surface tone patterns on Mende words, and unifies the behavior of nouns in isolation and nouns affixed with toneless suffixes.

However, such a map is not A-ISL. Jardine (2017b) shows that, surprisingly, a tone-map function like this is not even *monadic-second order* MSO definable. MSO logic is a generalization of the QF logic we explore here; in MSO, one can quantify not only over individual variables, but also over variables interpreted as sets of positions. Jardine (2017b) proves that this one-to-one indexing of tones to TBUs is beyond the expressivity of MSO. As MSO is a strict extension of QF, this implies that this tone map is also not QF.

Koser et al. (to appear) propose a solution that partially captures a notion of output locality. They invoke the notion of *least-fixed point* (LFP) logics, which allow for recursive definitions of the logical formulas that make up our transductions. The full notation of LFP is technical, so following Koser et al. we use *implicit definitions* (Rogers, 1998); that is, definitions that are explicitly recursive. For example, the formula for the left-to-right, one-to-one map in Mende is as below in (87).

$$(87) \quad A'(x, y) \stackrel{\text{def}}{=} \underbrace{((H(x) \vee L(x)) \wedge \text{first}(x) \wedge \sigma(y) \wedge \text{first}(y))}_{\text{a.}} \vee \underbrace{A'(p(x), p(y))}_{\text{b.}}$$

The disjunct in (87a) is the base case, or starting point of the recursion. It identifies the first tone in the melody and first syllable on the timing tier. The disjunct in (87b) is the recursive case: it states that $A'(x, y)$ is true of x and y when $A'(x, y)$ is true for the *predecessors* of both x and y . We can visualize the evaluation of this definition as depicted in the example in (88).




In (88) dashed lines represent pairs that are being evaluated at each step; solid lines show pairs that already have been evaluated. The first L tone and the initial syllable satisfy $A'(x, y)$ without any recursion; they satisfy the first disjunct (87a). Once we know the first L and the initial syllable satisfy $A'(x, y)$, then by (87b) we know their successors do: the H and the second syllable satisfy $A'(p(x), p(y))$ because the first L (the predecessor of H) and the initial syllable (the predecessor of the second syllable) satisfy $A'(x, y)$. Likewise, once we know that the H and the second syllable satisfy $A'(x, y)$, then by (87b) again we know

that the second L and the third syllable also satisfy $A'(x, y)$. In this way, the recursive definition of association maps tones to syllables much in the same way it was originally defined in the work of Leben (1973), Williams (1976), and Goldsmith (1976).

Unbounded spreading can be defined in the same way. In unbounded spreading, recursion starts at any underlying association.

$$(89) \quad A'(x, y) \stackrel{\text{def}}{=} \underbrace{A(x, y)}_{\text{a.}} \vee \underbrace{(A'(x, p(y)) \wedge \neg \text{final}(y))}_{\text{b.}}$$



In (90), the H and the initial syllable are associated in the input; thus, they satisfy the first disjunct (89a) $A(x, y)$ of the definition of $A'(x, y)$ in (89). Thus, this pair satisfies $A'(x, y)$. The recursive cases can then be evaluated. Because the initial syllable is the predecessor of the second syllable, and because the second syllable is not final, the H and the *second* syllable satisfy the recursive disjunct (89b). For the same reason, then, the H and the third syllable also satisfy this disjunct. In this way, the recursive nature of the definition captures the iterative addition of association lines. This recursion ends at the fourth syllable, which fails $\neg \text{final}(y)$, and so cannot satisfy the recursive disjunct (89b). Thus, the definition in (89) captures unbounded spreading to the penult.

The LFP definition in (89) of $A'(x, y)$ captures a notion of output locality, as it specifies that a pair are associated in the output if x and the predecessor of y are associated in the output. However, LFP logics are quite powerful, and thus we say that LFP logics only *partially* capture a notion of output locality.

For example, recursive definitions can also capture patterns that are neither input- nor output-local. Returning to string representations, consider the following recursively defined predicate that identifies any syllable that follows an underlying H-toned syllable.

$$(91) \quad \text{follows-}\acute{\sigma}(x) \stackrel{\text{def}}{=} \underbrace{\acute{\sigma}(x)}_{\text{a.}} \vee \underbrace{\text{follows-}\acute{\sigma}(p(x))}_{\text{b.}}$$

An example of (91) being evaluated over a string of syllables is given in Fig. 4. As with the above definitions of output association, this predicate can be thought of as being evaluated iteratively. However, note that it distinguishes the syllables that follow an H-toned syllable (i.e., those that are shaded in the last line of Fig. 4), from those that don't (those without any shading).





Figure 4: Recursive evaluation of $\text{follows-}\acute{\sigma}(x)$ in (91). Positions that satisfy the predicate at each step are highlighted in gray.

We can then use this to give a logical definition for unbounded shift, as witnessed by (92).

$$(92) \quad \acute{\sigma}'(x) \stackrel{\text{def}}{=} \text{penult}(x) \wedge \text{follows-}\acute{\sigma}(x)$$

The example below in (93) shows how this accurately distinguishes a penultimate syllable following an underlying H-toned syllable from one that does not.



$$\begin{array}{c}
\text{b.} \quad / \sigma \sigma \sigma \sigma \boxed{\sigma} \sigma / \\
\quad \quad [\sigma \sigma \sigma \sigma \boxed{\sigma} \sigma]
\end{array}$$

In (93a), the highlighted syllables show positions that satisfy $\text{follows-}\acute{\sigma}(x)$, parallel to Fig. 4. The penultimate syllable, highlighted with a box, also satisfies $\text{penult}(x)$, and so it satisfies the definition in (92) for $\acute{\sigma}'(x)$. This is contrasted with the penult in (93b), which satisfies $\text{penult}(x)$ but not $\text{follows-}\acute{\sigma}(x)$, and thus fails $\acute{\sigma}'(x)$ and is not output as $\acute{\sigma}$. Thus, the definition in (92) of $\acute{\sigma}'(x)$ outputs a syllable as high-toned if and only if it is a penult that follows a high tone. Unbounded shift, then, is a LFP-definable string function, whereas it was shown above to be neither ISL nor OSL.

As such, while LFP is promising as a logical characterization of phonological transformations, it does not strictly adhere to any notion of locality. A full exploration of the relationship of LFP to locality is best left to future work. However, it bears emphasizing that, with ARs, unbounded shift (and unbounded Meussen’s rule) can be captured *without* the recursive definitions of LFP. That is, the fact that ARs reduce these processes to simple, local characterizations will not change, regardless of what is revealed by future investigations of logical characterizations of phonology.

6 Conclusion

Computational analyses of phonological processes give us exact notions of locality that tell us precisely when a process is and is not local. In turn, logical transformations give us representation-independent definitions of computational locality, which allow us to directly compare different kinds of representation. This paper applied these techniques to study, in a rigorous way, the ability of ARs to describe long-distance processes in a local manner. This reveals that ARs do accomplish this in some cases—specifically, those in which changes on a tier are based on input information local to that change on that tier, or when association between two elements is determined by input information local to each element on their respective tiers.

However, there are cases—specifically, when changes are based on output information—that are not captured by considering ARs alone. Thus, while ARs do indeed contribute to a local conception of phonology, output locality is also necessary. The results here point to a full characterization of tone that includes both ARs and input- and output-based definitions of computational locality.

References

- Ao, B. (1991). Kikongo nasal harmony and context-sensitive underspecification. *Linguistic Inquiry*, 22(2):193–196.
- Bickmore, L. S. (1996). Bantu tone spreading and displacement as alignment and minimal misalignment. ROA #161-1196.
- Bickmore, L. S. and Kula, N. C. (2013). Ternary spreading and the OCP in Copperbelt Bemba. *Studies in African Linguistics*, 42.
- Cassimjee, F. and Kisseberth, C. (1998). Optimal domains theory and Bantu tonology. In Kisseberth, C. and Hyman, L., editors, *Theoretical Aspects of Bantu Tone*, pages 265–314. CSLI.
- Chandlee, J. (2014). *Strictly Local Phonological Processes*. PhD thesis, University of Delaware.
- Chandlee, J., Eyraud, R., and Heinz, J. (2014). Learning strictly local subsequential functions. *Transactions of the Association for Computational Linguistics*, 2:491–503.
- Chandlee, J., Eyraud, R., and Heinz, J. (2015a). Output strictly local functions. In Kornai, A. and Kuhlmann, M., editors, *Proceedings of the 14th Meeting on the Mathematics of Language (MoL 14)*, pages 52–63, Chicago, IL.
- Chandlee, J. and Heinz, J. (2018). Strictly locality and phonological maps. *LI*, 49:23–60.
- Chandlee, J. and Jardine, A. (2019a). Autosegmental input strictly local functions. *Transactions of the Association for Computational Linguistics*, 7:157–168.

- Chandlee, J. and Jardine, A. (2019b). Quantifier-free least fixed point functions for phonology. In *Proceedings of the 16th Meeting on the Mathematics of Language*, pages 50–62, Toronto, Canada. Association for Computational Linguistics.
- Chandlee, J., Jardine, A., and Heinz, J. (2015b). Learning repairs for marked structures. In *Proceedings of the 2015 Annual Meeting on Phonology*. LSA.
- Chandlee, J. and Lindell, S. (forthcoming). A logical characterization of strictly local functions. In Heinz, J., editor, *Doing Computational Phonology*. OUP.
- Chioşain, M. N. and Padgett, J. (2001). Markedness, segment realization, and locality in spreading. In Lombardi, L., editor, *Segmental phonology in Optimality Theory*, pages 118–156. Cambridge University Press.
- Clements, G. N. (1977). Neutral vowels in Hungarian vowel harmony: an autosegmental interpretation. In *NELS 7*, pages 49–64.
- Eisner, J. (1997a). Efficient generation in primitive Optimality Theory. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 313–320, Madrid.
- Eisner, J. (1997b). What constraints should OT allow? Talk handout, Linguistic Society of America, Chicago. ROA#204-0797. Available at <http://roa.rutgers.edu/>.
- Engelfriet, J. and Hooġboom, H. J. (2001). MSO definable string transductions and two-way finite-state transducers. *ACM Transactions on Computational Logic*, 2:216–254.
- Gafos, A. (1996). *The articulatory basis of locality in phonology*. PhD thesis, Johns Hopkins University.
- Goldsmith, J. (1976). *Autosegmental Phonology*. PhD thesis, Massachusetts Institute of Technology.
- Heinz, J., Rawal, C., and Tanner, H. G. (2011). Tier-based strictly local constraints for phonology. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 58–64, Portland, Oregon, USA. Association for Computational Linguistics.
- Hyman, L. (2011). Tone: Is it different? In Goldsmith, J. A., Riggle, J., and Yu, A. C. L., editors, *The Blackwell Handbook of Phonological Theory*, pages 197–238. Wiley-Blackwell.
- Hyman, L. (2014). How autosegmental is phonology? *The Linguistic Review*, 31:363–400.
- Hyman, L. and Katamba, F. X. (2010). Tone, syntax and prosodic domains in Luganda. In Downing, L., Rialland, A., Beltzung, J.-M., Manus, S., Patin, C., and Riedel, K., editors, *Papers from the Workshop on Bantu Relative Clauses*, volume 53 of *ZAS Papers in Linguistics*, pages 69–98. ZAS Berlin.
- Jardine, A. (2017a). The local nature of tone-association patterns. *Phonology*, 34:385–405.
- Jardine, A. (2017b). On the logical complexity of autosegmental representations. In Kanazawa, M., de Groote, P., and Sadrzadeh, M., editors, *Proceedings of the 15th Meeting on the Mathematics of Language*, pages 22–35, London, UK. Association for Computational Linguistics.
- Jardine, A. (2019). The expressivity of autosegmental grammars. *Journal of Logic, Language, and Information*, 28:9–54.
- Kager, R. (1999). *Optimality Theory*. Cambridge University Press.
- Kenstowicz, M. and Kisseberth, C. (1990). Chizigula tonology: the word and beyond. In Inkelas, S. and Zec, D., editors, *The Phonology–Syntax Connection*, pages 163–194. Chicago: the University of Chicago Press.
- Kisseberth, C. W. (1984). Digo tonology. In Clements, G. and Goldsmith, J. A., editors, *Autosegmental Studies in Bantu Tone*, pages 105–182. Foris Publications.
- Koser, N., Oakden, C., and Jardine, A. (to appear). Tone association and output locality in non-linear structures. In *Supplemental Proceedings of the 2019 Annual Meeting on Phonology*.
- Leben, W. R. (1973). *Suprasegmental phonology*. PhD thesis, Massachusetts Institute of Technology.

- Leben, W. R. (1978). The representation of tone. In Fromkin, V., editor, *Tone—A Linguistic Survey*, pages 177–219. Academic Press.
- Leben, W. R. (2006). Rethinking autosegmental phonology. In Mugane, J., editor, *Selected Proceedings of the 35th Annual Conference on African Linguistics*, pages 1–9. MA: Cascilla Proceedings Project.
- Levergood, B. (1987). *Topics in Arusa phonology and morphology*. PhD thesis, University of Texas.
- McCarthy, J. (1996). Remarks on phonological opacity in Optimality Theory. In *Studies in Afroasiatic Grammar: Papers from the Second Conference on Afroasiatic Linguistics, Sophia Antipolis, 1994*, pages 215–243. The Hague: Holland Academic Graphics.
- McCarthy, J. (2010). Autosegmental spreading in Optimality Theory. In Goldsmith, J. A., Hume, E., and Wezels, W. L., editors, *Tones and Features: Phonetic and Phonological Perspectives*, pages 195–222. De Gruyter Mouton.
- McCarthy, J. and Prince, A. (1993). Generalized alignment. In Booij, G. and van Marle, J., editors, *Yearbook of Morphology*, pages 79–153. Dordrecht: Kluwer.
- McCarthy, J. and Prince, A. (1995). Faithfulness and reduplicative identity. In Beckman, J., Dickey, L. W., and Urbanczyk, S., editors, *Papers in Optimality Theory*, number 18 in University of Massachusetts Occasional Papers in Linguistics, pages 249–384. University of Massachusetts.
- McCarthy, J. J. (2003). OT constraints are categorical. *Phonology*, 20:75–138.
- McMullin, K. and Hansson, G. O. (2016). Long-distance phonotactics as tier-based strictly 2-local languages. In *Proceedings of AMP 2015*. Linguistics Society of America.
- Myers, S. (1987). *Tone and the structure of words in Shona*. PhD thesis, University of Massachusetts, Amherst.
- Myers, S. (1997). OCP effects in Optimality Theory. *NLLT*, 15(4):847–892.
- Odden, D. (1982). Tonal phenomena in Kishambaa. *Studies in African Linguistics*, 13(2):177–208.
- Odden, D. (1986). On the role of the Obligatory Contour Principle in phonological theory. *Language*, 62(2):353–383.
- Odden, D. (1994). Adjacency parameters in phonology. *Language*, 70(2):289–330.
- Odden, D. (1995). Tone: African languages. In Goldsmith, J., editor, *The handbook of phonological theory*, pages 444–475. Oxford: Blackwell.
- Odden, D. (2001). Tone shift and spread in Taita I. *Studies in African Linguistics*, 30(1):76–110.
- Prince, A. and Smolensky, P. (1993). Optimality Theory: Constraint interaction in generative grammar. *Rutgers University Center for Cognitive Science Technical Report, 2*.
- Prince, A. and Smolensky, P. (2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell Publishing.
- Rogers, J. (1998). *A Descriptive Approach to Language-Theoretic Complexity*. Chicago: Center for the Study of Language and Information.
- Rogers, J. and Pullum, G. (2011). Aural pattern recognition experiments and the subregular hierarchy. *Journal of Logic, Language and Information*, 20:329–342.
- Schadeberg, T. (1979). Über die töne der verbalen formen im Rimi. *Afrika und Übersee*, 57:288–313.
- Shih, S. and Inkelas, S. (2019). Autosegmental aims in surface optimizing phonology. *Linguistic Inquiry*, 50(1):137–196.
- Sibanda, G. (2004). *Verbal phonology and morphology of Ndebele*. PhD thesis, UC Berkeley.
- Williams, E. S. (1976). Underlying tone in Margi and Igbo. *Linguistic Inquiry*, 7(3):463–484.
- Yip, M. (2002). *Tone*. Cambridge University Press.
- Zoll, C. (2003). Optimal tone mapping. *LI*, 34(2):225–268.